

Machine Learning to Support the Optimization of Graphene-Based Materials for Electrochemical Devices

Thesis submitted to University College London

by

David Noriega Pérez

Submitted in part fulfilment of the requirements for the degree of

Doctor of Philosophy

2026

Supervisors:

Primary: Dr. Rhodri Jarvis

Secondary: Dr. Thomas Miller

Electrochemical Innovation Lab

Department of Chemical Engineering

University College London

Declaration

I, David Noriega Pérez, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, it has been indicated in the thesis.

Acknowledgements

I would first like to thank Rhodri Jervis, Diego Díaz and Dan Brett for their support, guidance and technical mentoring throughout these four and a half years. I am also grateful to my managers at ArcelorMittal and TheNextPangea, Nicolás de Abajo and Rodrigo García, for offering me this opportunity and for supporting and facilitating the process of making my PhD compatible with my professional role. My sincere thanks also go to my second supervisor, Thomas Miller, for his expert insights on graphene supercapacitors and the complexities of their characterisation and scale-up.

Thanks to Siva and Roberto for fostering the relationship between UCL Electrochemical Innovation Lab and ArcelorMittal Global R&D Avilés and for encouraging me to embark on this adventure.

I would also like to thank my colleagues and friends at the laboratory in Avilés for their help with synthesis and characterisation tasks, especially to Carlota, Cristina, Tan, Marcos, Luis Miguel and Guillermo. Many thanks to my colleagues in London for welcoming me so warmly during my visits and for their support with the lab work, especially Zahra, Hanieh and Srinivas.

Finally, heartfelt thanks to my partner, my friends from Valderas, Asturias, Madrid and Berlin, and my family, for their love and unwavering support. You gave me the strength to carry out this work.

Abstract

The transition to a sustainable, zero-emissions society demands the discovery of new materials that are scalable and cost-effective. Traditional materials development is often slow and inefficient. This thesis develops a physics-informed Machine Learning (ML) framework to accelerate the synthesis optimisation and performance prediction of graphene-based materials for electrochemical energy applications.

A literature review on the role of iron in electrochemical systems motivated the exploration of Kish graphite, a byproduct of the steel industry, as sustainable and iron-containing carbon source. This choice aligns with the goal of circular materials design by transforming industrial waste into high-value catalysts and supports. Kish-derived materials were investigated as catalysts for the Oxygen Reduction Reaction (ORR) and as a support for platinum in fuel cells. The high sensitivity of Rotating Disk Electrode (RDE) measurements to experimental noise, such as variations in ink drying or electrode preparation, compromised reproducibility. In contrast, capacitance evaluations using Constant Current Charge-Discharge (CCCD) tests on rGO synthesized from commercial graphene oxide (GO) demonstrated more consistent and reliable results. This motivated a focused effort on optimizing capacitance as a more data-efficient target for ML-guided synthesis.

A physics-informed Gaussian Process model was developed to optimise the chemical reduction of GO using ascorbic acid. The model incorporates experimentally derived noise and enforces physically meaningful constraints to ensure scalability. Bayesian Optimisation was used to efficiently identify optimal synthesis conditions, with the model validated across different experimental setups and literature data. In the final stage, a ML model specific to rGO-based supercapacitors was constructed. The model focused exclusively on graphene-derived materials measured via CCCD in KOH electrolyte, ensuring experimental consistency but increasing the risk of overfitting. This was mitigated through domain-informed steps including Mahalanobis outlier detection, Leave-One-Out cross-validation, and physical interpretability checks. Key drivers of capacitance identified include nitrogen and oxygen doping and average pore width.

Overall, this work demonstrates a reproducible and physically grounded ML methodology that integrates experiment-guided synthesis optimisation with data-driven performance modelling, establishing a transferable framework for sustainable materials discovery in electrochemical systems.

Impact Statement

This industrial thesis, conducted in collaboration with ArcelorMittal Global R&D, demonstrates how ML, when combined with domain knowledge and experimental data, can accelerate the discovery of new materials development for energy applications under constraints of limited data and physical resources. It addresses key challenges in experimental materials science, such as experimental noise, limited scalability, and the tendency of ML models to overfit when applied to specific material families.

The importance of this work is reinforced by recent events, such as the blackout in Spain, which revealed the vulnerabilities of current energy infrastructures and the urgent need for more resilient energy systems. In this context, the thesis contributes to the development of advanced materials for supercapacitors, which enhance grid stability.

A central strength of this work is its industrial relevance. Kish graphite, a byproduct of the steel industry, is investigated as a sustainable carbon source for synthesising reduced graphene oxide (rGO). This aligns with circular economy goals and demonstrates how industrial waste materials can be repurposed for high-performance applications in energy storage.

Natural graphite is exposed to price fluctuations driven by EV-battery demand and are classified as critical raw materials by the European Commission. However, the availability of Kish is linked to the blast-furnace route, which is expected to decline as the steel sector decarbonises. The approach demonstrated here should therefore be viewed as a medium-term opportunity to add value to an existing waste stream, rather than a permanent solution to graphite supply constraints.

Environmental sustainability also extends to the chemistry employed, and the thesis considers the potential role of bio-derived reductants such as ascorbic acid as greener alternatives to conventional reagents. In addition, it presents a critical analysis of the potential for iron to move beyond its conventional structural role and emerge as a functional component in electrochemical devices.

A major impact of this work lies in its methodological innovation. The strategic prioritisation of galvanostatic charge-discharge tests, based on their higher reproducibility compared to rotating disk electrode measurements, demonstrates how uncertainty analysis can shape scientific strategy and improve research efficiency. Moreover, the developed Gaussian Process models incorporate experimentally measured noise and physically motivated constraints, improving robustness, interpretability, and generalisability.

Additionally, the experimental boundaries applied during the chemical reduction process were defined by principles of scalability and compatibility with industrial processing. The application of Bayesian Optimisation to explore synthesis parameters reduced the need for extensive experimentation, saving time and resources.

Finally, the thesis develops a carefully curated ML model that incorporates physical correctness, addressing the issue of overfitting in small, material-specific datasets. Through publication-level cross-validation, targeted feature selection, and physical interpretability checks, the model identifies nitrogen and oxygen contents and pore structure as consistent drivers of capacitance.

The approaches developed here are directly applicable to the broader energy materials community, including work on battery electrodes, electrocatalysts, membranes, and sustainable polymers. More broadly, this thesis contributes to the foundation of reproducible, physics-aware ML in materials science, an essential step for moving from lab-scale insights to industrial-scale impact.

Table of Contents

| | | |
|--------|--|----|
| 1 | Introduction | 18 |
| 1.1 | Motivation..... | 18 |
| 1.2 | Literature Review: The Role of Iron in Electrochemical Devices | 23 |
| 1.2.1 | Iron Basics..... | 24 |
| 1.2.2 | Iron in PEMFC..... | 30 |
| 1.2.3 | Iron in Lithium-Ion Batteries..... | 32 |
| 1.2.4 | Iron in Lithium-Sulphur Batteries..... | 35 |
| 1.2.5 | Iron in Redox Flow Batteries..... | 37 |
| 1.2.6 | Super-Iron Batteries | 41 |
| 1.2.7 | Iron-Air Batteries..... | 43 |
| 1.2.8 | Nickel-Iron Batteries | 45 |
| 1.2.9 | Other Uses | 47 |
| 1.2.10 | Research Gap | 50 |
| 1.3 | Literature Review: The Role of Graphene in Supercapacitors | 52 |
| 1.3.1 | Machine Learning for the Design of Experiments..... | 57 |
| 1.3.2 | Gaussian Processes and Bayesian Optimisation in Graphene Research | 60 |
| 1.3.3 | Machine Learning Models and Explainability of Graphene Capacitance | 65 |
| 1.3.4 | Research Gap | 77 |
| 2 | Methods and Materials | 79 |
| 2.1 | Contribution Statement..... | 79 |
| 2.2 | Experimental Techniques | 79 |
| 2.2.1 | Graphene Synthesis | 79 |
| 2.2.2 | Graphene Oxide and Reduced Graphene Oxide from Kish Graphite | 80 |
| 2.2.3 | Expansion of Kish Graphite | 82 |
| 2.2.4 | Cyclic Voltammetry..... | 83 |
| 2.2.5 | Linear Sweep Voltammetry and Rotating Disk Electrode | 85 |
| 2.2.6 | Constant Current Charge/Discharge..... | 86 |
| 2.2.7 | Raman Spectroscopy | 89 |
| 2.2.8 | SEM | 91 |
| 2.2.9 | Elemental Analysis | 92 |
| 2.2.10 | Gas Physisorption..... | 92 |
| 2.3 | Machine Learning Techniques..... | 94 |
| 2.3.1 | Parametric vs Non-Parametric Models | 94 |
| 2.3.2 | Gaussian Processes..... | 96 |
| 2.3.3 | Bayesian Optimisation..... | 99 |

| | | |
|-------|--|-----|
| 3 | Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide | 105 |
| | Summary | 105 |
| 3.1 | Introduction | 105 |
| 3.2 | Results | 112 |
| 3.2.1 | Experimental | 112 |
| 3.2.2 | Galvanostatic Charge-Discharge Tests | 121 |
| 3.3 | Conclusions and Future Work | 127 |
| 4 | Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes | 129 |
| | Summary | 129 |
| 4.1 | Introduction | 129 |
| 4.2 | Results | 133 |
| 4.2.1 | Formal Definition of the Gaussian Process Regression Model | 133 |
| 4.2.2 | Uncertainty Quantification and Model Constraints | 134 |
| 4.2.3 | Surrogate Modelling and Experimental Optimisation | 142 |
| 4.2.4 | Interlaboratory Modelling | 149 |
| 4.3 | Conclusions and Future Work | 155 |
| 5 | Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity | 157 |
| | Summary | 157 |
| 5.1 | Introduction | 157 |
| 5.2 | Methods | 159 |
| 5.2.1 | Feature Selection | 159 |
| 5.2.2 | Data Collection | 161 |
| 5.2.3 | Data Preprocessing | 161 |
| 5.2.4 | Model Selection | 165 |
| 5.3 | Results | 167 |
| 5.3.1 | Physical Correctness Assessment | 173 |
| 5.3.2 | Comparative Evaluation of GP and RF Predictions | 183 |
| 5.4 | Conclusions and Future Work | 187 |
| 6 | Conclusions and Future Work | 190 |
| 6.1 | Conclusions | 190 |
| 6.1.1 | Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide | 191 |
| 6.1.2 | Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes | 191 |
| 6.1.3 | Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity | 192 |

| | | |
|-------|---|-----|
| 6.2 | Future Work | 193 |
| 6.2.1 | Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide | 193 |
| 6.2.2 | Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes | 193 |
| 6.2.3 | Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity..... | 194 |
| | APPENDIX A..... | 197 |
| A.1 | Variability in RDE Polarisation Curves for Pt40%/Vulcan | 197 |
| A.2 | Discharge Time Variability at 1 A/g and 10 A/g Under Different Reduction and Stirring Conditions..... | 198 |
| | APPENDIX B..... | 199 |
| B.1 | Experimental Estimation of the Irreducible Noise | 199 |
| B.2 | Derivation of $I_s = 2.3$ from an Assumption of an Inter-Setup Tolerance $\leq 15\%$ | 200 |
| | APPENDIX C | 202 |
| C.1 | Estimation of Average Pore Diameter from Published Pore Size Distributions..... | 202 |
| C.2 | Publication Cross-Referencing for Dataset Analysis | 203 |
| C.3 | Distributions of Input Features and Capacitance..... | 204 |
| | REFERENCES..... | 208 |

List of Abbreviations

| | |
|--------|-------------------------------------|
| ANN | Artificial Neural Network |
| BET | Brunauer–Emmett–Teller |
| BJH | Barrett-Joyner-Halenda |
| BO | Bayesian Optimisation |
| BRR | Bayesian Ridge Regression |
| BTRL | Battery Technology Readiness Level |
| CCCD | Constant Current Charge/Discharge |
| CNTs | Carbon Nanotubes |
| DFT | Density Functional Theory |
| DoE | Design Of Experiments |
| DTR | decision tree regression |
| EDLC | Electric Double Layer Capacitor |
| EI | Expected Improvement |
| ESR | Equivalent Series Resistance |
| G | Graphene |
| GC | Gas Chromatography |
| GLR | Generalized Linear Regression |
| GO | Graphene Oxide |
| HER | Hydrogen Evolution Reaction |
| HOR | Hydrogen Oxidation Reaction |
| HTGO | Hydrothermal Porous Graphene |
| ICRFB | Iron Chromium Redox Flow Battery |
| KNN | K-Nearest Neighbours |
| LFP | LiFePO ₄ |
| LHS | Latin Hypercube Sampling |
| LSV | Linear Sweep Voltammetry |
| LTC | Lattice Thermal Conductivity |
| ML | Machine Learning |
| MOF | Metal Oxide Framework |
| NN | Neural Networks |
| OEMs | Original Equipment Manufacturers |
| ORR | Oxygen Reduction Reaction |
| PEMFCs | Proton-Exchange Membrane Fuel Cells |
| PI | Probability of Improvement |
| RDE | Rotating Disk Electrode |
| RF | Random forest |
| RFB | Redox Flow Battery |
| RGO | Reduced Graphene Oxide |
| RMB | Rechargeable Magnesium Batteries |
| RPD | Robust Parameter Design |
| RSM | Response Surface Methodology |
| SCs | Supercapacitors |
| SE | Sequential Experimentation |
| SEM | Scanning Electron Microscope |
| SHAP | SHapley Additive exPlanations |

| | |
|-----|-----------------------------------|
| SVM | Support Vector Machine |
| TDP | Temperature Programmed Desorption |
| TGA | Thermal Gravimetric Analysis |
| TRL | Technology Readiness Level |
| UCB | Upper Confidence Bound |
| XRD | X-Ray Diffraction |

List of Tables

| | |
|---|-----|
| Table 1-1. Examples of iron complexes as active materials in fuel cells and redox flow batteries..... | 26 |
| Table 1-2. Summary of studies on lithium-ion batteries using iron-based materials. | 34 |
| Table 1-3. Overview of research on lithium-sulphur batteries with iron-based components..... | 37 |
| Table 1-4. Summary of studies on iron redox flow batteries. | 41 |
| Table 1-5. Overview of iron's roles in various energy storage technologies, including the specific form used, associated technical challenges, technological maturity, and key references..... | 50 |
| Table 1-6. Summary of specific capacitance of different graphene materials (partially extracted from [214–216])..... | 56 |
| Table 1-7. Summary of ML models predicting capacitance based on characteristics of carbonaceous material, including graphene..... | 75 |
| Table 2-1. Open Source Software Libraries for Bayesian Optimisation. | 102 |
| Table 3-1. Composition of Reduced Graphene Oxide from Kish graphite. | 112 |
| Table 3-2. Deposition and drying conditions for different catalyst inks on top of working electrode prior to ORR activity evaluation. | 114 |
| Table 3-3. Electrochemical parameters from polarisation curves for Expanded Kish Graphite, Pt20%/rGO and Pt40%/Vulcan samples..... | 115 |
| Table 3-4. Electrochemical parameters from polarisation curves for Pt40%/Vulcan replicates for different ink concentrations. *This Coef. Variation was only calculated for the pair (0.73 mg/mL, 4 drops) and (2.94 mg/mL, 1 drop). | 117 |
| Table 3-5. Summary of KPIs extracted from ORR polarisation curves of Pt40%/Vulcan catalyst films measured under two experimental setups. | 120 |
| Table 3-6. Discharge times (in seconds) at different applied voltages (0, -0.4, -0.8, -1.2 V) for rGO-based electrodes tested at 2 A/g, fabricated under three reduction conditions (65 °C, 3 h; 80 °C, 0.5 h; 95 °C, 0.5 h) using two stirring methods (magnetic and paddle). The coefficient of variation (CV, %) between stirring methods is shown for each condition, illustrating the reproducibility of electrode performance across processing variables. | 124 |
| Table 3-7. Summary of specific capacitance values calculated and CoVs for each condition. | 126 |
| Table 4-1. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for electrodes with rGO from reductions using two different heating procedures. The bottom row corresponds to the standard deviation between them. | 138 |
| Table 4-2. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes with different mass. The bottom row corresponds to the standard deviation between them. | 139 |
| Table 4-3. Calculation of the total noise for 1 A/g, 2 A/g and 10 A/g..... | 140 |
| Table 4-4. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes produced under three different reduction conditions: 80°C for 5 hours, 65°C for 2 hours and 85°C for 1.75 hours. | 143 |
| Table 4-5. Characteristics of the two set-ups used to evaluate the capacitance of rGO electrodes. | 150 |
| Table 4-6. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes measured in two different setups. | 151 |
| Table 4-7. Noise attributed to measurements coming from different setups, for 1A/g, 2A/g and 10A/g. | 151 |
| Table 4-8. Total noise considered for 1 A/g, 2 A/g and 10 A/g. | 152 |

| | |
|--|-----|
| Table 5-1. Technical relevance and presence in previous model of the selected physicochemical features of graphene. | 159 |
| Table 5-2. Summary of performance metrics for the five models during validation and testing | 169 |
| Table 5-3. Physicochemical characterisation of selected rGO samples. | 184 |

List of Figures

| | |
|---|----|
| Figure 1-1. Battery Technology Readiness Level (BTRL). Reproduced from Crabtree [5]. | 20 |
| Figure 1-2. Summary of the various roles of iron in electrochemical devices. | 24 |
| Figure 1-3. Pourbaix Diagram of Iron in water at 25°C. Reproduced under CC BY-SA 3.0 license from Metallos [22]. | 27 |
| Figure 1-4. Iron-carbon phase diagram, reproduced from Föll [43]. | 29 |
| Figure 1-5. (a) Activity and performance evolution of PGM-free catalysts from 2010 to 2020. Full symbols: not iR-corrected; hollow symbols: iR-corrected. Synthesis method: CB, carbon support + organic precursor; HT, hard template; MOF, metal organic framework. Reproduced with permission from Osmieri et al. [56]. (b) Volcano Plots for Two- and Four-Electron ORR Pathways: DFT Onset Potentials and Experimental Current Densities as a Function of HO* Binding Free Energy (GHO). Reproduced with permission from Sun et al. [68]. | 32 |
| Figure 1-6. Overview of joining tasks in battery applications: schematic depiction of the joining location (a) for cylindrical cells and (b) for pouch or prismatic cells; (c) battery module consisting of cylindrical cells directly connected by a large busbar, reproduced from [153]; (d) single pouch cell interconnected to a copper collector bar, reproduced from [154] with permission from Elsevier. The figure is reproduced from [34] under the CC BY-NC-ND license. | 48 |
| Figure 1-7. Graphene markets by application by 2025, reproduced with permission from Reiss <i>et al.</i> [194]. | 53 |
| Figure 1-8. Efficiency of the Bayesian optimisation. $N_{\text{Bayes/rand}}$ is the average number of calculations needed until, for the first time, a structure that belongs to the top k% of all the candidates is found with the Bayesian/random search. Reproduced with permissions from [263]. | 61 |
| Figure 1-9. SHAP plots of ML models for graphene capacitance: (left) [295], (right) [298], reproduced with permission. | 77 |
| Figure 2-1. Collection spots of Kish graphite from the ArcelorMittal factory in Avilés, Spain. | 81 |
| Figure 2-2. Visual progression of material transformation. Left: Raw Kish graphite sample. Centre: Aqueous suspension of graphene oxide obtained from oxidised Kish graphite. Right: Dried graphene oxide powder after exfoliation and purification. | 82 |
| Figure 2-3. Suspension of ascorbic acid with reduced graphene oxide at the bottom. | 82 |
| Figure 2-4. Steps for the expansion of Kish Graphite using chemical and microwave processes. | 83 |
| Figure 2-5. Left: Schematic representation of an electrochemical cell for CV experiments. Reproduced from [330]. Right: Cyclic voltammogram of polycrystalline Pt (1 1 1) in 0.5 M HClO ₄ at 10 mV/s, from [331]. | 84 |
| Figure 2-6. Illustration of CCD test result from Electric Double Layer Capacitors or Pseudocapacitors with linear potential change over time. Reproduced with permission from [334]. | 87 |
| Figure 2-7. Schematic diagrams and equivalent circuits for (a) a three-electrode setup, (b) a two-electrode setup with equal electrode masses, and (c) a two-electrode setup with different electrode masses. Reproduced with permission from [334]. | 88 |
| Figure 2-8. Methodology for the preparation of rGO electrodes for CCD tests. | 89 |
| Figure 2-9. Dependence of the Raman spectra on number of layers and disorder [336]. Left: Raman spectra of graphite, metallic and semiconducting nanotubes, low and high sp ³ amorphous carbons; Right: 2D peak as a function of N for 514 nm excitation. Flakes with N layers are indicated by NLG. Reproduced with permission. | 90 |

| | |
|---|-----|
| Figure 2-10. The Raman spectra of graphite (top), GO (middle), and the reduced GO (bottom). Reproduced with permission from [338]. | 91 |
| Figure 2-11. Prior and posterior distribution using Gaussian processes with radial basis function kernel after 9 observations [351]. | 99 |
| Figure 2-12. Description of an active learning cycle. Reproduced with permission from [255]. | 103 |
| Figure 3-1. Left: distribution of potential for each of the laboratories (indicated by different colours) needed to reach 2, 5 and 10 mA·cm ⁻² . Right: potential at 10 mA·cm ⁻² before (CV 1) and after (CV 2) the stability measurement. Reproduced from Ref. [359] with permission from the Royal Society of Chemistry. | 107 |
| Figure 3-2. Photographs of different catalyst films deposited on glassy carbon electrodes, prepared by the same researcher. Ref. [359] with permission from the Royal Society of Chemistry. | 108 |
| Figure 3-3. Comparison of best and worst practices for the electrochemical Characterisation of supercapacitors (Table 1, top), and reproducibility obtained using the best practices across different potentiostats (Table 2). Reproduced with permission from [360]. | 109 |
| Figure 3-4. Top: SEM images of reduced graphene oxide powder; bottom: Raman spectra, with an I _D /I _G ratio typical from RGO powders, showing an intense D band related to a high presence of defects. | 113 |
| Figure 3-5. Polarisation curves for RGO20%Pt (red and purple lines), Vulcan40%Pt (orange and green lines) and Expanded Kish (blue line). | 115 |
| Figure 3-6. Left: photographs of catalyst films deposited on glassy carbon electrode under different preparation conditions. Right: Polarisation curves of Pt40%/Vulcan corresponding to a catalyst loading of 30 mg Pt·cm ⁻² at different deposition conditions. | 117 |
| Figure 3-7. Experimental setups: Setup 1 (left) and Setup 2 (right). | 119 |
| Figure 3-8. ORR polarisation curves for Pt40%/Vulcan catalyst films measured under two experimental setups. | 120 |
| Figure 3-9. Swagelok cell used to extract galvanostatic charge-discharge (GCD) curves with a three-electrode setup. | 123 |
| Figure 3-10. GCD curves under different reducing conditions and mechanical stirring method, measured at current densities of 1 A/g (left), 2 A/g (centre), and 10 A/g (right). | 124 |
| Figure 3-11. Specific capacitance values calculated using Equation 3-4 for all tested conditions. | 125 |
| Figure 4-1 Conceptual scheme illustrating the main sources of experimental uncertainty affecting the measured specific capacitance. The contributions considered include mass variability during electrode preparation, temperature fluctuations during graphene oxide reduction and measurement-protocol variability. In a reduction process based on the chemical route at fixed temperature, a typical lab procedure consists of introducing the reactants in a flask placed on a hot plate, as illustrated in left photo of Figure 4-2. The increase in the temperature within the flask is limited by connecting it to a condenser, whose cold reservoir is an ice bath. This helps maintain a constant water volume during the reduction. | 135 |
| Figure 4-2. Experimental set-ups for the reduction of graphene oxide using ascorbic acid as reducing agent under controlled temperature. Left: set-up with temperature control via a hot plate. Middle: setup with temperature control via an Arduino-based system, which measures bath temperature using a thermocouple and adjusts power to an electric blanket accordingly. Right: temperature and time curves recorded for both set-ups during an identical reduction treatment at 80°C for 30 mins. | 136 |

Figure 4-3. GCD curves of rGO electrodes reduced using two different heating systems. The curves compare electrodes whose graphene oxide was reduced at 80°C for 30 mins using either a conventional hotplate (dashed line) or an Arduino-controlled heating blanket (solid lines). Galvanostatic charge-discharge tests were performed at 1 A/g (left), 2 A/g (middle) and 10 A/g (right)..... 138

Figure 4-4. GCD curves for reduced graphene oxide electrodes. The graphene oxide was reduced at 80°C for 30 mins. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). The dashed lines correspond to electrodes containing **2.46 mg** of rGO, while the solid line are electrodes with **2.75 mg** of rGO. 139

Figure 4-5. Graphene production pilot plant. Photo courtesy of TheNextPangea®. ... 142

Figure 4-6. GCD curves for reduced graphene oxide electrodes under three different reduction conditions. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). The dashed lines correspond to 80°C and 0.5 hours, the solid line to 65°C and 3 hours, and the dotted line to 95°C and 0.5 hours. 143

Figure 4-7. Gaussian Processes after three capacitance measurements at 2 A/g, with Temperature (x-axis) and Time (y-axis). The crosses on the plots indicate the points of real data (observations). The top plot shows the predicted mean distribution. The bottom plot illustrates the uncertainty in terms of the 95% confidence interval, calculated as 1.96 times the standard deviation σn^2 , where σn^2 is computed according to Equation 2-7. 145

Figure 4-8. Expected Improvement for the Gaussian Process estimated from capacitance values of graphene oxide electrodes reduced at 80°C (0.5h), 65°C (3h) and 95°C (0.5h). The red star indicates the point where the EI reaches its maximum value in the first iteration. 145

Figure 4-9. GP mean prediction distribution (top) and 95% confidence level (bottom) after 13 observations. 147

Figure 4-10. Residual plot for each held-out condition during LOO validation. 147

Figure 4-11. Top: Nyquist plots for rGO electrodes reduced under different temperature and time conditions, measured at 0.8V (solid lines) 1.2V (dash-dotted lines). Bottom: corresponding CV curves for the same electrodes and voltage conditions. 149

Figure 4-12. GCD curves for two different reduced graphene oxide electrodes measured in two different experimental set-ups in two different labs. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). Top figures correspond to measurements carried out using set-up 1, bottom figures refer to measures using set-up 2. The graphene in the electrodes was reduced from graphene oxide at 65°C for 2 hours. 151

Figure 4-13. Residual plot for each held-out condition during LOO validation. Black dots represent capacitance values measured with Setup-1, while red dots correspond to Setup-2. The blue dot indicates the residual of the test point, whose reference value was taken from ref. [423]. The prediction for this point was computed using the following inputs: 95°C, 1h and Setup-1. 153

Figure 4-14. Residual plot for each held-out condition during LOO validation under the two setup configurations. 154

Figure 5-1. Left: Mahalanobis distances across publications from which datapoints were extracted to build the dataset. Right: Deviation of each publication's mean Mahalanobis distance. 164

Figure 5-2. Real versus predicted capacitance values for each testing fold under Leave-One-Publication-Out Cross-Validation. The dispersion reflects both model error and inter-publication variability inherent to literature-derived experimental data. The

| | |
|--|-----|
| figure illustrates consistency of predictive behaviour across heterogeneous studies rather than pointwise prediction accuracy. | 168 |
| Figure 5-3. Real versus predicted capacitance values for the Random Forest model evaluated using two different random 80-20% train-test splits of the full dataset. Although the visual distributions appear similar, the resulting performance metrics differ, reflecting sensitivity to data partitioning. The improved apparent performance relative to Figure 5-2 arises from the presence of samples from the same publications in both training and test sets, leading to optimistic estimates of predictive accuracy. | 171 |
| Figure 5-4. Summary plot of SHAP values demonstrating the influence of features on the Random Forest Regressor model. | 172 |
| Figure 5-5. Mean Absolute Differences between consecutive cumulative means of the SHAP values for all the features after 100 iterations. | 173 |
| Figure 5-6. Distributions of Average pore width (top) and Pore volume (bottom) with key metrics. | 176 |
| Figure 5-7. SHAP Dependence plot for average pore width, with colours referring to pore volume values. | 176 |
| Figure 5-8. SHAP Dependence Plot for Pore volume vs N% (left) and O% (right). | 177 |
| Figure 5-9. SHAP Dependence Plot for O% vs N%. | 178 |
| Figure 5-10. SHAP dependence plot of SSA, with colour indicating oxygen values. | 179 |
| Figure 5-11. Spearman correlation matrix showing relationships between model features and capacitance. | 181 |
| Figure 5-12. SHAP Dependence Plots for I_D/I_G ratio vs. N, O, and S content (top, middle and bottom, respectively). | 182 |
| Figure 5-13. SHAP dependence plot of S%, with colour indicating Specific Surface Area values. | 183 |
| Figure 5-14. Comparison between experimentally measured capacitance values (Setups 1 and 2) and predictions obtained from the Gaussian Process (GP) and Random Forest (RF) models for selected rGO samples described in Table 5-3. Symbols (●, ■, ▲, ✖, ⊕) uniquely identify individual samples listed in Table 5-3 and are used consistently across the table and figure. Error bars on the GP predictions represent the 95% confidence intervals. The dashed line indicates ideal agreement between predicted and measured capacitance. | 185 |
| Figure 5-15. Left: Mahalanobis distances across publications from which datapoints were extracted to build the dataset, together with the datapoints from Table 5-3 corresponding to the experimental work described in Chapter 4. Right: Deviation of each publication's mean Mahalanobis distance. | 186 |
| Figure 5-16. Representative SEM micrographs of the circle and triangle samples at three magnifications. (a-c) Circle sample. (d-f) Triangle sample. | 187 |

1 Introduction

1.1 Motivation

One of the most pressing challenges in the energy transition is the rapid development of electrochemical devices that utilize active materials with low environmental impact. A common strategy to address this involves replacing highly CO₂-intensive materials with alternatives that require less energy and emit fewer greenhouse gases during production. This approach is currently driving global materials research, with the goal of achieving electrochemical performance comparable to that of commercially available materials.

However, the targets set by international agreements and government agencies are highly ambitious. The Paris Climate Agreement, adopted in 2015, aims to limit global warming to well below 2°C, with efforts to restrict it further to 1.5°C. Achieving these goals necessitates reaching carbon neutrality by 2050 [1,2]. At the same time, technical targets for hydrogen storage systems and fuel cell devices, as defined by the U.S. Department of Energy [3], require levels of electrochemical performance that current materials, whether already available or nearing market readiness, fail to meet.

Research into new materials for electrochemical devices faces several key challenges:

- Long and technically complex research and development cycles
- High initial capital investment required to scale up production of newly developed active materials
- Limited funding opportunities for high-risk, potentially disruptive research

These challenges have hindered the emergence of economically viable material technologies for batteries, particularly among new companies in Europe and the United States. As a result, recent progress has primarily come from incremental improvements made by large corporations with the resources and risk tolerance to invest heavily in research and development. Most of these enhancements focus on increasing lithium-ion cell capacity and are led by multinational chemical companies.

Nevertheless, there is a growing consensus that incremental improvements alone will be insufficient to meet the future energy demands of a zero-emissions economy, particularly in terms of both energy density and power density.

One key reason for the limited progress in battery innovation is the intense pressure that original equipment manufacturers (OEMs) place on battery suppliers, which leaves little room for disruptive innovation. Additionally, in an era dominated by demanding shareholders and large-scale mergers (such as the Dow–DuPont merger), market pressure to deliver short-term profits often results in reduced R&D budgets. This environment favours low-risk, incremental improvements and discourages impactful and potentially transformative advances in energy storage technologies.

Commercialising new materials, whether for batteries or other applications, is notoriously difficult. This holds true for both large corporations and start-ups. A published analysis [4] highlights the stark reality for battery technology start-ups: among 36 companies analysed, only two achieved positive exits, meaning that the return on investment exceeded the capital invested (return > 1×). Seventeen of the remaining companies managed to reach early commercialisation, defined as either pilot-scale manufacturing or the establishment of a development agreement with a major industry player. Even in these relatively successful cases, the average time required to reach early commercialisation was eight years.

Several key factors explain the long timelines, high risks, and substantial costs associated with innovation in electrochemical energy storage technologies, particularly lithium-ion, sodium-ion, and next-generation battery systems (e.g., solid-state and multivalent batteries):

- The synthesis of new materials typically begins at the laboratory scale, with small batch sizes and mostly manual procedures. Scalability is not a primary concern at this stage, as the focus is on discovering novel properties.
- Achieving precise control over materials at the nanoscale is key. The targeted performance of new storage systems depends on carefully engineered physicochemical properties. This necessitates rigorous control over synthesis methods, extensive and accurate characterisation of the materials, and high batch-to-batch reproducibility.

As a result, the development cycles for electrochemical energy storage materials and battery technologies often span 10 to 15 years from conception to realisation in real applications, as an energy storage device must simultaneously satisfy challenging requirements in terms of energy density, power capability, safety, lifetime, and cost under realistic operating conditions. The initial stages are usually supported by university-based research, with scale-up studies typically beginning around the fifth year. It may then take an additional five years to achieve full techno-economic validation. Figure 1-1

below illustrates the concept of Battery Technology Readiness Levels (BTRL), which provides a structured framework for assessing the maturity of battery technologies from early-stage research to full industrial deployment. The figure maps the different development phases, ranging from fundamental materials discovery at laboratory scale to large-scale manufacturing and market implementation, highlighting the increasing complexity, validation requirements and investment associated with each stage.

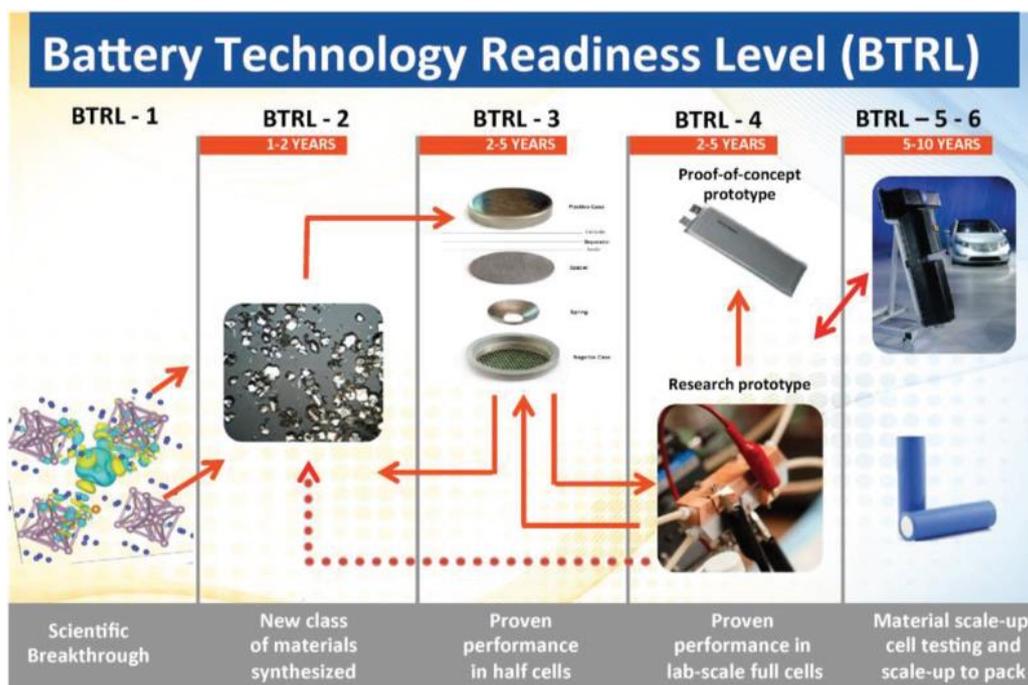


Figure 1-1. Battery Technology Readiness Level (BTRL). Reproduced from Crabtree [5].

The figure shows that the pathway from laboratory innovation to commercial deployment in electrochemical energy-storage systems is long and resource-intensive, typically requiring more than a decade of sustained effort. This reflects the difficulty of simultaneously optimising electrochemical performance, safety, durability and cost under realistic operating conditions.

Over the past three decades, advances in computational science have greatly contributed to materials discovery, leading to the development of numerous computationally designed composites and the on-demand availability of *ab initio* predicted properties.

Currently, the most widely used approaches for understanding and predicting materials synthesis combine experimental observation, physics-based modelling, and data-driven learning. In situ X-Ray Diffraction [6] (XRD) provides real-time experimental insight into phase transformations during synthesis. *Ab initio* thermodynamic modelling [7,8], grounded in quantum-mechanical calculations, captures the fundamental energetic and

kinetic factors governing material formation. In parallel, Machine Learning (ML) methods [9,10] have emerged as a complementary approach, capable of identifying empirical relationships and optimising synthesis parameters when direct physical modelling is computationally intractable or when data from multiple sources must be integrated. In organic chemistry, ML-assisted synthesis design has benefited from extensive reaction databases such as Reaxys [11], which contains more than 57 million elementary reactions. These reactions enable the use of retrosynthetic analysis, a method that breaks down target molecules into simpler precursor structures, in order to design scalable synthesis routes based on commercially available starting materials.

Although there are also several databases for the *ab initio* calculation of electronic structures in inorganic materials (Materials Project [12], Open Quantum Materials Database [13], Novel Materials Discovery [14], AFLOW [15]), these primarily contain calculations of perfect crystalline structures, two-dimensional materials and/or gas-phase molecules. As a result, a retrosynthetic-like analysis for inorganic compounds is not possible. The prediction of synthesis methods for these materials remains limited, as the existing data does not capture the complexities of real-world synthesis.

In this context, the application of ML to inorganic synthesis is typically focused on identifying empirical patterns in fabrication techniques. This is achieved through large-scale text mining of the scientific literature using natural language processing techniques, ranging from classical NLP methods to modern large language models (LLMs), rather than through first-principles prediction of reaction pathways.

Most recent advances involve developing natural language processing (NLP) pipelines to mine synthesis recipes, precursor selections, and reaction conditions from published papers, enabling the construction of large, structured datasets for downstream ML tasks [16]. However, it is important to note that while the primary focus is on empirical pattern recognition via text mining, there is growing interest in integrating these data-driven methods with physics-based or first-principles approaches [17].

Currently, there is no database that systematically compiles analysable information from the thousands of experiments conducted daily in universities and research centres focused on the synthesis of new materials. This lack of structured and accessible data makes it difficult, and sometimes even impossible, to transfer computational predictions to the laboratory or to reproduce the fabrication of promising materials.

ArcelorMittal is one of the world's leading steel companies. It holds a dominant position in major global steel markets, including automotive, construction, household appliances, packaging, and energy [18]. The company is also one of the world's leading producers

of iron ore, with more than 40 million tonnes of iron produced in 2024. Decarbonisation targets of ArcelorMittal are the most ambitious of any steelmaker, with the objective of reducing CO₂ emissions up to 35% by 2030 and reaching net-zero by 2050. These efforts are supported by multi-billion-dollar investments in the company operations, which are aimed to ensure technological readiness towards the future zero-emissions steelmaking.

To support the company's net-zero strategy, the following two research questions have been formulated and are addressed in the literature review of this thesis:

1. How can iron, in its various forms, contribute to the development of new electrochemical devices?
2. How can ML support research on graphene as a material for supercapacitors?
3. How can physics-informed ML approaches integrate insights from both iron- and carbon-based systems to accelerate the discovery of sustainable electrochemical materials?

The first question is addressed by initially explaining the fundamental electrochemical properties of iron, followed by a discussion of how these properties translate into functional roles in electrochemical devices, supported by several examples. The second is explored through a review of the current state of graphene-based supercapacitors, including recent commercial developments, and key findings from the scientific literature, with particular attention to the role of ML.

Based on insights drawn from these literature reviews, two research gaps are identified and described. These gaps form the foundation for the chapters presenting the results of the investigation conducted in this thesis; in chapter 3 Kish, a steelmaking byproduct containing graphite, is analysed as a catalyst for the Oxygen Reduction Reaction (ORR). The experimental uncertainty associated with this analysis is described and compared with that of an analogous study on the capacitance of rGO chemically reduced under different time and temperature conditions. Based on the findings of this chapter, Chapter 4 explores the use of physics-informed Gaussian Processes, an increasingly adopted ML technique, to guide the synthesis of rGO materials with the aim of maximising their capacitance. Finally, Chapter 5 introduces a new ML model that relates rGO capacitance to its fundamental physicochemical features. Innovative techniques are applied to both the model construction and evaluation stages, aiming to address the limitations of current models identified in the literature review.

1.2 Literature Review: The Role of Iron in Electrochemical Devices

Iron is one of the most abundant elements on Earth. The development of human technology has been strongly linked to this element since the earliest days of our history, serving as disruptive material for advanced weapons and tools in the Iron Age (~1200 BC), supporting the development of high-distance transportation of iron-locomotives running on iron-rails that led the industrial revolution, and providing an outstanding and tremendously versatile offer of steels for automotive that allowed the expansion of this industry in the 20th century.

Steel, an alloy primarily composed of iron and carbon, represents the most economically and technologically important form of iron today. More than 3,500 different grades [19] of steels are available with many different physical, chemical, and environment properties, out of which around 75% have been developed in the last 20 years. Modern industry is now about to pass through a new radical transformation, with the ambitious target of replacing the energy-intensive model that was established in the previous industrial revolutions by a new one with no CO₂ emissions and with a much more sustainable use of materials.

Being 100% recyclable, iron and steel are clearly seen as examples of materials that fit quite well in this circular economy picture, so construction and mobility sectors will still require steel mills as a source of material. What is not so obvious is the notorious importance that iron is gaining in the field of new electrochemical devices and systems. This section of the literature review aims to present a general and critical view about the role of iron in electrochemical devices. Electrochemical devices are systems that convert or store energy through chemical reactions occurring at the interface between electrodes and electrolytes. They include technologies such as fuel cells, batteries, and supercapacitors, all of which rely on materials capable of efficiently catalysing or storing charge through redox or interfacial processes. Starting from some basic properties of iron from an electrochemical point of view, a suite of references including key actors in the energy storage field of research are introduced. Finally, other uses of iron in the future green economy are described.

Figure 1-2 provides an overview of the diverse functions that iron and iron-based materials can play in different electrochemical energy-conversion and storage systems. The schematic groups the roles of iron according to device type, including its use as an active catalytic centre or current collector for fuel cells, as a cathode or anode in batteries, and as a structural material for battery casing or hydrogen pipelines.

The figure emphasises the versatility of iron in electrochemical systems, ranging from its role in molecular catalysis to its use as a structural material in device construction. This multi-level functionality forms the conceptual basis for the literature review that follows.

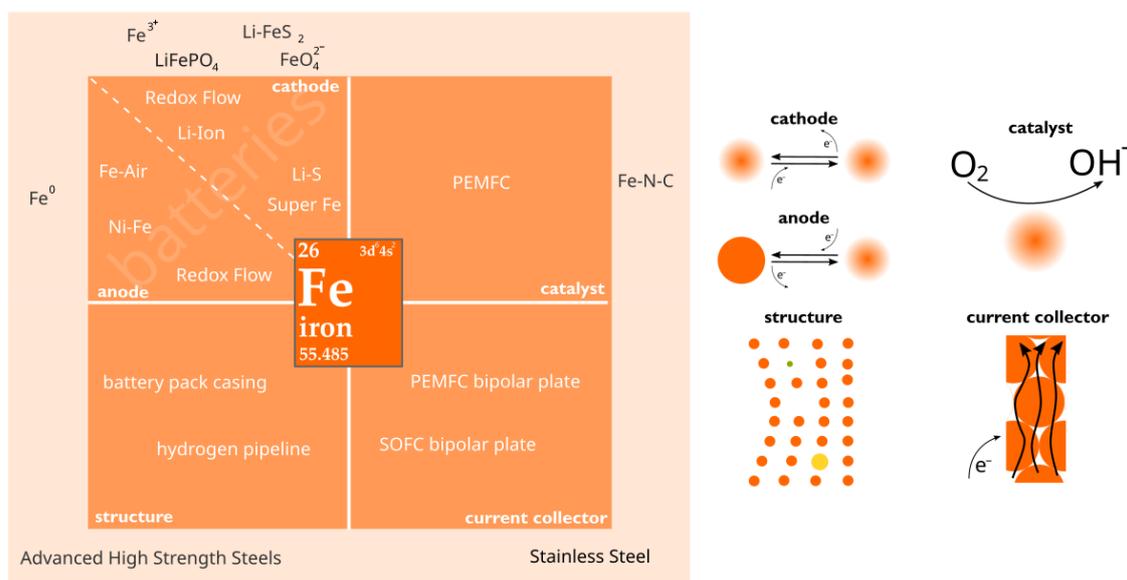


Figure 1-2. Summary of the various roles of iron in electrochemical devices.

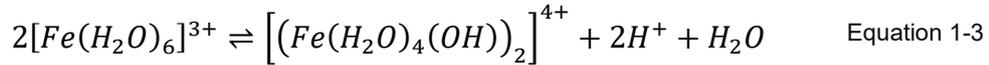
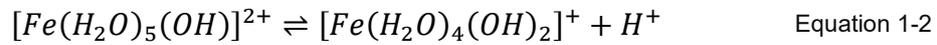
1.2.1 Iron Basics

Iron is a metal with atomic number 26 and 55.845 standard atomic weight. In its solid form, iron can form crystals with different arrangements named allotropes and denoted by the Greek letters α , γ , δ and ϵ . Of those, α -Iron is the stable phase below 912°C. In this phase, iron has a body-centred cubic crystal structure, and it is ferromagnetic up to 767.85°C, known as its Curie temperature. Pure iron is a silvery-white soft metal of little use in industry. This is because it rusts rapidly in moist air and its mechanical properties are limited compared to carbon-alloyed counterparts.

The electron configuration of iron is $[\text{Ar}] 3d^6 4s^2$, and its known oxidation states range from -II (d^{10}) to +VIII (d^0). Of those, the most stable and widespread are II (d^6) and III (d^5) corresponding to the ferrous and ferric states respectively [20].

Iron (II) exists in various coordination geometries and spin states, forming salts with nearly all anions. This ability to create a wide range of complexes is key for understanding its opportunities and limitations as an electrochemically active material. Complexing agents can change the reduction potential of iron (II) in aqueous solution as well as its stability. The majority of iron (II) complexes are octahedral; however, tetrahedral and five-coordinate compounds are also known [21].

Iron (III) forms salts with all common anions except for those which are strong reducing agents such as iodide. In the absence of coordinating ligands, the following equilibria are important:



If the pH rises above 2-3, colloidal gels begin to form, eventually leading to the precipitation of reddish-brown hydrous iron (III) oxide. This precipitate is responsible for the low solubility of iron (III) at pH 7. In the presence of complexing anions like chloride, the hydrolysis of iron (III) becomes more complex, resulting in the formation of chloro, aqua, and hydroxy species, as well as $[FeCl_4]^-$ at high chloride concentrations.

Due to its versatile chemistry, iron complexes have been studied thoroughly as active materials in numerous works related to fuel cells and redox flow batteries. Table 1-1 below summarises some of these works.

Figure 1-3 shows the Pourbaix diagram of iron in aqueous solution at 25 °C. The diagram maps the thermodynamically stable phases of iron as a function of pH and electrochemical potential, indicating the regions where metallic iron, iron oxides/hydroxides, or soluble iron ions are expected to be stable. The water-stability window is also shown, delimiting the domain where hydrogen or oxygen evolution reactions are not thermodynamically favoured [22].

In the figure above, the region of water stability is shown as a dotted blue line. The boundaries of the stability domains, meaning the areas in the Pourbaix diagram where each phase of iron (metallic, oxide, or hydroxide) is stable against corrosion or transformation, are as follows:

1. $Fe_{(s)} \rightleftharpoons Fe^{2+} + 2e^- \quad || \quad E_0 = -0.44 + 0.0295 \log Fe^{2+}$
2. $Fe^{2+} \rightleftharpoons Fe^{3+} + e^- \quad || \quad E_0 = 0.771 + 0.0591 \log \frac{Fe^{3+}}{Fe^{2+}}$
3. $2Fe^{3+} + 3H_2O \rightleftharpoons Fe_2O_{3(s)} + 6H^+ \quad || \quad E_0 = -0.72 - 3pH$
4. $2Fe^{2+} + 3H_2O \rightleftharpoons Fe_2O_{3(s)} + 6H^+ + 2e^- \quad || \quad E_0 = 0.728 - 0.1773pH - 0.0591 \log Fe^{2+}$
5. $3Fe^{2+} + 4H_2O \rightleftharpoons Fe_3O_{4(s)} + 8H^+ + 2e^- \quad || \quad E_0 = 0.980 - 0.2364pH - 0.0886 \log Fe^{2+}$
6. $2Fe_3O_{4(s)} + H_2O \rightleftharpoons 3Fe_2O_{3(s)} + 2H^+ + 2e^- \quad || \quad E_0 = 0.221 - 0.0591pH$

Table 1-1. Examples of iron complexes as active materials in fuel cells and redox flow batteries.

| Complex | Role | Type of reaction | Solvent | Ref. |
|---|-----------|------------------|---|------|
| N-doped porous carbon that anchors both atomically dispersed Fe-N ₄ sites and Fe atomic clusters (FeAC@FeSA-N-C) | Catalyst | ORR | Water + KOH, methanol | [23] |
| graphene encapsulated Fe/Fe ₃ C nanocrystals-Fe-N _x configurations (Fe@C-FeNC) | | | HClO ₄ with addition of NaSCN, KOH | [24] |
| [2Fe ₂ S] complex | | HER | 1:1 DMF:acetate | [25] |
| Fe-phosphido dimer complex | | | THF: H ₂ O | [26] |
| [FeFe]-hydrogenase attached to a fluorene or silafluorene sensitizer | | | aqueous MeCN | [27] |
| [FeFe]-hydrogenase linked to pyridylphosphole encapsulated in the caged complex PSCageFe ₄ (Zn-L) ₆ , formed of six zinc porphyrin ligands linked by four iron(II) centers, | | | CH ₃ CN, TFA or H ₂ O | [28] |
| tris(2,2'-bipyridine)iron(II)tetrafluoroborate | Catholyte | Fe oxidation | propylenecarbonate | [29] |
| iron(III) acetylacetonate | | | 84/16 vol.% acetonitrile/1,3-dioxolane | [30] |
| [Fe(phen) ₃] ^{2+/3+} (phen = 1,10-phenanthroline) | | | CH ₃ CN | [31] |
| Fe(CN) ₆ | | | Water + KOH | [32] |
| Fe(DIPSO) | Anolyte | Fe reduction | Water + KOH | [32] |
| [tris(imino)pyridine]Fe][OTf] ₂ | | | acetonitrile | [33] |

According to Figure 1-3, iron can be classified as a base metal, as its thermodynamic stability domain at 25 °C and atmospheric pressure does not overlap with that of water. Consequently, iron is unstable in the presence of water and many aqueous non-oxidising solutions, and it corrodes in such environments with the evolution of hydrogen gas. This reaction is particularly intense in acidic solutions and gradually decreases in intensity as the pH increases. Between approximately pH 10 and 13, iron does not corrode due to the formation of a protective oxide film on its surface. However, at pH values above 13, even in the absence of oxidising agents, the solutions become corrosive once again.

The electrode potential of iron immersed in solution increases in the presence of oxidising agents. The addition of such compounds can either lead to partial or complete passivation of the metal, or, conversely, enhance its corrosion rate. The outcome

depends on whether the resulting increase in potential is sufficient to shift the metal into the passivation domain.

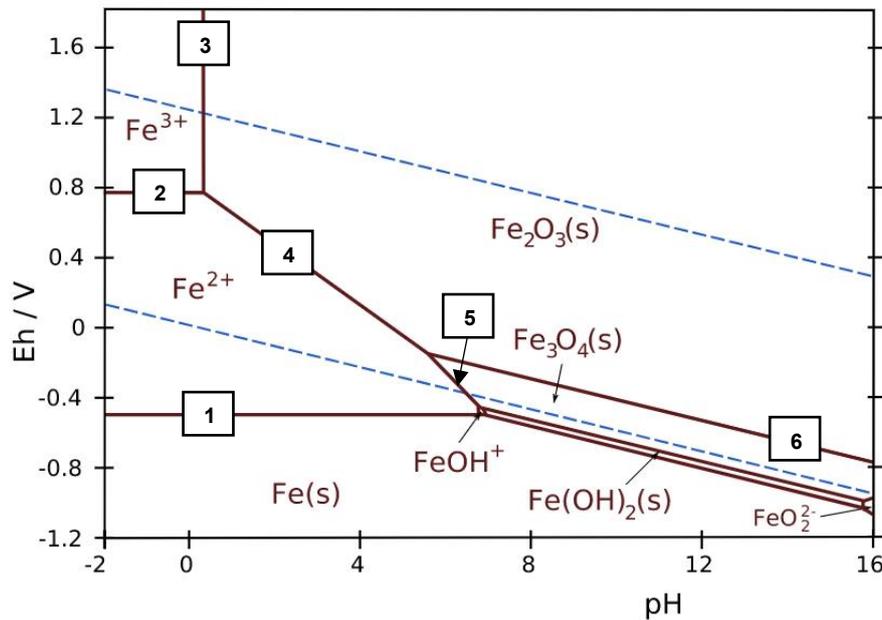


Figure 1-3. Pourbaix Diagram of Iron in water at 25°C. Reproduced under CC BY-SA 3.0 license from *Metallos* [22].

The ease with which iron is passivated by oxidising agents increases as the range of electrode potentials allowing corrosion becomes narrower. Passivation in iron is relatively difficult, or even impossible, at pH values below approximately 8. It becomes progressively easier at pH values above 8 and is particularly favourable in the range between pH 10 and 12.

Steel is an alloy made up of iron with a few tenths of a percent of carbon and with a mix of different iron-family crystalline domains, filled with vacancies, dislocations, interstitial and substitutional atoms coming from different alloying elements like manganese, nickel, chromium and many others. Steel is the most widely used metal in the world and the second most consumed material by humans, after concrete. The combination of strength, toughness, weldability, electrical conductivity, magnetic properties and recyclability makes steel the material of choice for multiple massive applications like automotive, energy structures, construction, packaging and house appliances.

The unique combination of steel's outstanding properties is also being considered in the design of structures for both energy storage and generation, serving as a key material to accelerate the energy transition to a net zero carbon emission economy. Steel is present in different parts of Electric Vehicles' battery pack system [34], electric motors [35,36],

wind turbines and solar installations [37], fuel cell bipolar plates [38], metal-air and redox flow batteries, to name a few examples.

Steel can be produced via two primary routes: the electric arc furnace route and the blast furnace route. The electric arc furnace (EAF) process uses scrap metal as the raw material, which is melted by the high temperatures generated from an electric arc formed between the scrap and graphite electrodes inside the furnace. In contrast, the blast furnace route begins with iron ores as the primary input. These ores are smelted in the blast furnace to produce liquid iron, which is subsequently refined in the steel shop to produce molten steel.

During the transfer of liquid iron from the blast furnace (typically around 1400-1500°C) to the steel shop, it is common for the temperature to drop by approximately 100 °C [39]. While this temperature loss may be viewed as a drawback in terms of energy efficiency, it also presents a secondary effect that could be considered an opportunity. Specifically, this effect can be explained by the temperature-dependent solubility of carbon in liquid iron, as described in the Iron-Carbon phase diagram (Figure 1-4, below). This diagram describes the equilibrium phases formed as a function of temperature and carbon content. It defines the phase fields corresponding to austenite, ferrite, cementite and liquid iron, as well as the eutectic and eutectoid transformations that occur during cooling.

As mentioned before, the temperature of iron when it is extracted from the blast furnace is around 1400-1500°C and it has a carbon percentage between 4 and 4.5% [39]. Thus, according to Figure 1-4, this means that a temperature loss under these conditions generates an excess in carbon which crystallises as graphite. In fact, this graphite also contains iron oxides, magnesium oxides and other oxides and silicates, either on top or intercalated between the graphene layers, and it is known as Kish graphite [40–42]. The origin of the oxide/silicate impurities is the use of scrap and other feed materials to both mitigate the loss of temperature and the oxidation of iron during its transportation.

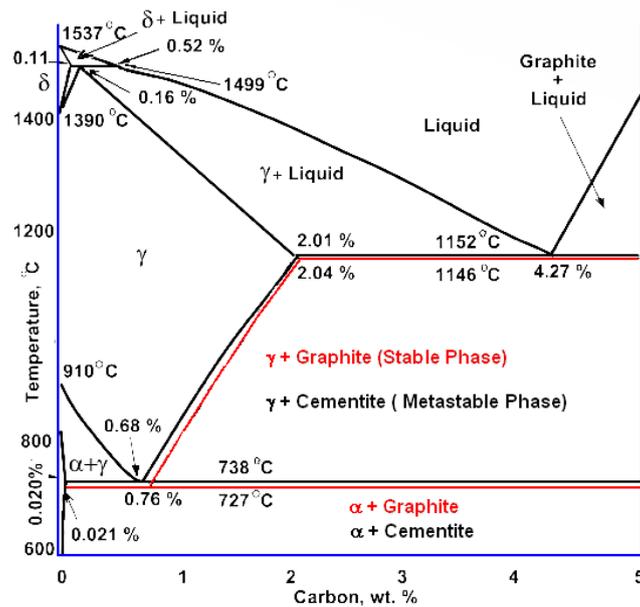


Figure 1-4. Iron-carbon phase diagram, reproduced from Föll [43].

Kish graphite is a well-known material in the graphene research community, and it is used by hundreds of research groups to make graphene by mechanical exfoliation, pioneered by the Nobel Laureates [44]. Kish graphite is an excellent precursor for graphene as predicted by several studies [45,46]. Graphene coming from Kish Graphite has five times higher electron mobility than pure chemical vapour deposition (CVD) made graphene, according to Japanese researchers [47]. Another study from Korean researchers showed that graphene coming from Kish has less defective structure than when it is obtained from Natural Graphite [48]. The origin for the higher crystallinity may be the formation of the crystal under liquid iron, a very smooth surface that energetically favours the generation of large basal crystal domains. Kish has also been studied as cathode for aluminium chloride-graphite batteries [49], achieving energy and power densities up to 65 Wh·kg⁻¹ and 4363 W·kg⁻¹, respectively.

Given the electrochemical versatility of iron and its complexes, multiple examples of research work that use iron as active element for energy storage can be found both in the literature and in the market [50,51]. Moreover, as previously discussed, Kish graphite, as a byproduct of iron manufacture, can be used as a source of pure graphite and its multiple nanosized forms like nanographene, multilayer graphene or graphene, which can be potentially used as electrodes and/or conductive additives in multiple electrochemical devices. A summary of some of these activities is presented here, with the focus on the use of iron in Proton-Exchange Membrane Fuel Cells (PEMFCs), Li-ion, Li-Sulphur and redox-flow batteries, followed by the potential of graphene, which can be extracted from Kish, as active material in supercapacitors.

1.2.2 Iron in PEMFC

Proton Exchange Membrane Fuel Cells (PEMFCs) are electrochemical devices that convert the chemical energy of hydrogen and oxygen directly into electricity, with water and heat as the only by-products. Nanosized platinum on carbon (Pt/C) is the most commonly used catalyst for PEMFCs. However, the high cost, insufficient activity and limited stability of commercial Pt/C make it difficult to meet the US DOE targets [52]. Among these targets, the United States Department of Energy cost goals for integrated PEMFC power systems and fuel cell stacks operating on direct hydrogen for transportation applications is 40\$·kW⁻¹ at 500,000 systems per year production [53]. With platinum as reference material, cost breakdown estimations for 0.125 mg_{Pt}·cm⁻² systems indicate that catalyst would be responsible for more than 40% of the total cost at 500,000 systems per year production rate [54], i.e. the largest single component of the PEMFC stack cost at high volume production. Thus, there is a motivation to reduce or eliminate platinum levels in catalysts to both meet cost targets and mitigate dependence on precious metal imports.

Compared with noble metal catalysts, non-noble metal catalysts, such as transition metal–nitrogen–carbon (M–N–C, M = Fe, Co, Ni, Cu, etc.) catalysts or transition metal oxide catalysts, have natural electrocatalysis and price advantages [55]. In a progressive evolution from 2010 to 2020, synthesis approaches diverged from those centred around carbonaceous supports and hard templates to those targeting carbon-rich matrices generated from M-, N-, and C-containing metal organic framework (MOF) precursors [56].

As previously described, iron complexes have been studied thoroughly as active materials in numerous works related to fuel cells and redox flow batteries due to their versatile chemistry. Iron-series elements, due to their unique electronic structure and low crystal field splitting energy, are facile for forming metal complexes with weak field ligands (such as ligands with P and N as coordination atoms) [57]. High-spin iron series complexes can also change the spin state through spin crossover, resulting in the reduction of the reaction energy barrier for promoting the Oxygen Reduction Reaction (ORR). As a result, iron is one of the most preferred candidate metals to replace precious metals as ORR catalysts. In recent years, iron series compounds, such as oxides, chalcogenides, carbides, phosphates and nitrides, have been incorporated within carbon matrix as high-performance ORR catalysts [58,59]. The introduction of carbon matrix not only provides an excellent conductivity framework but also prevents the agglomeration of iron-series compounds [60]. This agglomeration prevention enhances the presence of single-atom sites, with higher activity than clusters and nanoparticles [61]. At the same

time, iron-series elements can modify the electronegativity of carbon, thereby enhancing the catalytic activity of the material [62].

Dominguez *et al.* [63] studied the performance of Fe, Co, and Mn as non-precious metal catalysts (NPMCs) for the oxygen reduction reaction (ORR) in acidic conditions. They use a 4-step process with physicochemical, thermal, and acid treatments to prepare NPMCs starting from M-phthalocyanines and N-modified active carbon as precursors. They find that pyrolysis is essential for creating active and stable catalysts by incorporating the metal and forming graphitic-N groups. Fe-based NPMCs show higher ORR activity than Co and Mn-based ones, indicating that the choice of transition metal is crucial, with Fe being the most effective.

The superior properties of Fe as a catalyst are also supported by ab-initio calculations. Luo *et al.* [64] combined Density Functional Theory (DFT) calculations with Machine Learning techniques to identify high-performance carbon-based double-atom/nanocluster electrocatalysts for ORR under alkaline conditions. A total of 1200 potential catalysts were screened, including three metal-nitrogen-co-doped graphene coordination structures and 20 different elements, covering transition metals, noble metals, and rare-earth elements. Guided by ML, the authors found that the Fe-Ce double-atom catalyst is the best performing ORR catalyst in the sample space. In their recent study, Wu and Zelenay [65] analysed the current understanding of how the atomic structure of M-N-C catalysts correlates with their activity and stability under acidic conditions, supported by DFT calculations.

Stability under low-pH conditions is in fact one of the major challenges for the realisation of commercial Fe-N-C based catalyst that can replace Pt in PEMFCs [65,66]. As stated by Holby *et al.* [67], demetallation is thermodynamically favourable over FeN₄C₁₃₈ phases under acidic conditions, limiting the long-term stability of active sites in PEMFC cathodes. Figure 1-5 illustrates the superior catalytic activity of iron-based complexes compared to other alternatives that are also free of platinum group metals (PGM-free).

In this recent study [68], the authors present an innovative method to enhance the catalytic performance of Fe-based single-atom catalysts using a salt-template polymerisation technique. This method adjusts both the first and second coordination shells of Fe-N₄ isolated atoms. According to the authors, this dual adjustment redistributes charge density, lowers the d-band centre, and weakens the binding of oxygenated intermediates, thereby improving oxygen reduction reaction activity. A single Fe atom electrocatalyst was designed and synthesized with this optimized structure in

nitrogen and sulphur co-doped graphene. $\text{FeCl}_3 \cdot 6\text{H}_2\text{O}$ was used as the precursor, with Fe(III) acting as both initiator and oxidant for the in situ oxidative polymerisation of rhodamine monomers.

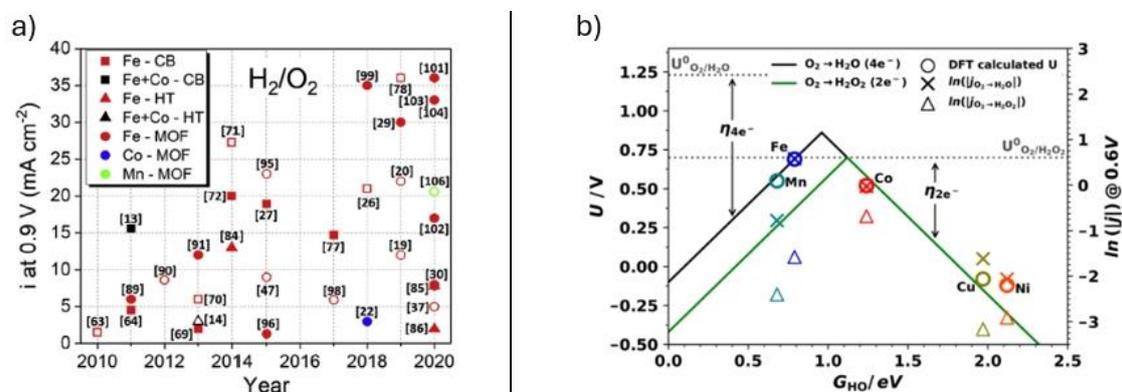


Figure 1-5. (a) Activity and performance evolution of PGM-free catalysts from 2010 to 2020. Full symbols: not iR-corrected; hollow symbols: iR-corrected. Synthesis method: CB, carbon support + organic precursor; HT, hard template; MOF, metal organic framework. Reproduced with permission from Osmieri et al. [56]. (b) Volcano Plots for Two- and Four-Electron ORR Pathways: DFT Onset Potentials and Experimental Current Densities as a Function of HO^* Binding Free Energy (G_{HO}). Reproduced with permission from Sun et al. [69].

Additionally, Fe acted as a catalyst during the subsequent pyrolysis of the complex, promoting the decomposition of carbon precursors and leading to the formation of graphene sheets. This catalyst surpasses benchmark Pt/C in ORR activity and shows superior stability in Zn-air batteries, paving the way for advanced single metal atom catalysts through precise atomic structure regulation.

1.2.3 Iron in Lithium-Ion Batteries

LiFePO_4 is a well-established cathode material for lithium-ion battery applications that has found a recent renaissance due to its low-cost and superior safety characteristics to high-energy density materials such as LiNiMnCoO_2 (NMC), and because of the relative safety of the material can be manufactured into large-format cells for electric vehicles, increasing cell-to-pack energy density. These properties also make them suitable for grid-scale energy storage devices for intermittent power generation by windmills and solar cells. Among the several candidates of choice for lithium-ion battery cathode materials, LiFePO_4 is reported to be inexpensive, nontoxic, and environmentally benign with a flat potential of 3.5 V versus lithium [70].

Significant efforts have been made to address the inherently low electrical conductivity of LiFePO_4 [71–74]. The most common approach to overcome this is by coating the LiFePO_4 nanoparticles with carbon. Carbon coatings are usually formed during LiFePO_4 synthesis, in which an organic precursor (the carbon source) and the inorganic raw

materials are mixed together. Then, heating the mixture in an inert or reducing atmosphere results in the simultaneous formation of conducting carbon and LiFePO_4 . Cai *et al.* [75] synthesized LFP cathode materials using amorphous carbon, carbon nanotubes (CNTs), and graphene (G) as conductive materials via sand milling and spray drying processes and followed by calcination. The initial discharge specific capacity increases in the order of LFP/C ($150.3 \text{ mAh}\cdot\text{g}^{-1}$) < LFP/C/CNTs ($155.7 \text{ mAh}\cdot\text{g}^{-1}$) < LFP/C/G ($159.7 \text{ mAh}\cdot\text{g}^{-1}$) < LFP/C/G/CNTs ($164.5 \text{ mAh}\cdot\text{g}^{-1}$) at 0.1 C. LFP/C/G/CNTs also exhibited considerable discharge specific capacity ($99.5 \text{ mAh}\cdot\text{g}^{-1}$) at high current of 5 C. The superior electrochemical performance of LFP/C/G/CNTs is attributed to the synergistic effect of CNTs and G on decreasing charge transfer impedance by constructing highly three-dimensional conductive network.

Three-dimensional (3D) porous LiFePO_4 /graphene aerogel (LFP/GA) composite was successfully prepared by an in-situ hydrothermal process by Du *et al.* [76]. An aerogel composite with 10% of graphene displays the best electrochemical performance, with the specific discharge capacities of $168 \text{ mAh}\cdot\text{g}^{-1}$ and $155 \text{ mAh}\cdot\text{g}^{-1}$ at 0.1C and 1C, respectively, and the capacity retains 96.3% for up to 800 cycles.

Chien *et al.* [77] tested graphene oxides (GOs) of various morphologies as conductive additives, including pristine GO, three-dimensional GO, and hydrothermal porous GO (HTGO) as a component of LiFePO_4 /carbon composite cathode using a method involving sol-gel processing, spray-drying, and calcination.

Ma *et al.* synthesized low-cost organometallic complex, $\text{Zn}(\text{OAc})_2$ -diethanolamine ($\text{Zn}(\text{OAc})_2\cdot\text{DEA}$) [78] as a binder and an additive in the LFP electrode. The researchers reported a high capacity of $169 \text{ mAh}\cdot\text{g}^{-1}$ for this material, approaching the theoretical capacity of $170 \text{ mAh}\cdot\text{g}^{-1}$, at the rate of 0.2C, $150 \text{ mAh}\cdot\text{g}^{-1}$ at the rate of 1C, and $86 \text{ mAh}\cdot\text{g}^{-1}$ at a high rate of 20C. The electrode also has excellent cycling performance, showing a low-capacity decay rate of 0.03% per cycle over 1500 cycles at 5C. These are relevant improvements in comparison to the nonmodified PVDF-based LFP electrode.

Literature on LiFePO_4 is predominantly based on nano-sized materials because this range is beneficial to the high-power applications as it can shorten the solid-state diffusion distance within LiFePO_4 . In addition, coating LiFePO_4 micro particles with uniform carbon coatings by calcination of carbon sources is difficult to achieve because inhomogeneous mixing between the carbon sources and the LiFePO_4 particles. In fact, a straightforward way to overcome the problem of low electronic conductivity is to reduce the particle size to the nanometer range. However, this size range limits the tap density, reducing the volumetric energy density of the battery. On the other hand, the calcination

process may be environmentally problematic. Guo *et al.* [71] described the utilisation of micron-sized LiFePO_4 , which has a higher tap density than its nano-sized counterparts, by applying a conducting polymer coating on its surface with a greener diazonium chemistry. The coated micron-sized LiFePO_4 , compared with its pristine counterpart, has shown improved electrical conductivity, high-rate capability and excellent cyclability when used as a 'carbon additive free' cathode material for rechargeable Li-ion batteries.

Table 1-2. Summary of studies on lithium-ion batteries using iron-based materials.

| Material | mAh/g | Rate (C) | Ref |
|--|-------|----------|------|
| LiFePO_4 (1 mm) + Polyphenylene | 165 | 0,1 | [71] |
| | ~135 | 1 | [71] |
| | ~115 | 10 | [71] |
| LiFePO_4 (10-200nm) + Amorphous Carbon | 150,3 | 0,1 | [75] |
| LiFePO_4 (10-200nm) + AC + CNTs | 155,7 | 0,1 | [75] |
| LiFePO_4 (10-200nm) + AC + rGO | 159,7 | 0,1 | [75] |
| LiFePO_4 (10-200nm) + AC + CNTs + rGO | 164,5 | 0,1 | [75] |
| LiFePO_4 (10-200nm) + AC + CNTs + rGO | ~145 | 1 | [75] |
| LiFePO_4 (10-200nm) + AC + CNTs + rGO | 100 | 5 | [75] |
| LiFePO_4 (2 mm) + rGO-5% | 159 | 0,1 | [76] |
| LiFePO_4 (1 mm) + rGO-10% | 168 | 0,1 | [76] |
| LiFePO_4 (1 mm) + rGO-10% | 155 | 1 | [76] |
| LiFePO_4 (1 mm) + rGO-10% | 131 | 5 | [76] |
| LiFePO_4 (1 mm) + rGO-10% | 115 | 10 | [76] |
| LiFePO_4 + rGO | 160,2 | 0,1 | [77] |
| | 132,2 | 10 | [77] |
| LiFePO_4 + $\text{Zn}(\text{OAc})_2 \cdot \text{DEA}$ | 169 | 0,2 | [78] |
| | 150 | 1 | [78] |
| | 86 | 20 | [78] |
| V-doped LiFePO_4 | 155,6 | 0,1 | [79] |
| | ~130 | 1 | [79] |
| | ~100 | 3 | [79] |

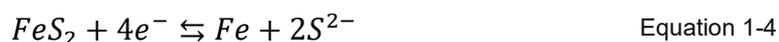
In this work [79], several V-doped LiFePO_4/C composites were prepared via a solid-state method. The $\text{LiFe}_{1-x}\text{VPO}_4/\text{C}$ ($x=0, 0.01, 0.02, 0.03$) cathode materials (V0, V1, V2, V3) exhibited enhanced electrochemical performance. V2 samples showed discharge capacities as high as 155.6, 150.2, 136.0, 122.4 and 93.1 $\text{mAh}\cdot\text{g}^{-1}$ at 0.1, 0.2, 0.5, 1.0 and 3.0C, respectively, and approximately 100% capacity retention after 100 cycles at 1.0C. Table 1-2 provides a summary of recent studies on lithium-ion batteries utilising iron-based materials, highlighting their compositions, synthesis methods, and electrochemical performance.

1.2.4 Iron in Lithium-Sulphur Batteries

Lithium-Sulphur batteries have theoretical specific capacities as high as $1675 \text{ mAh}\cdot\text{g}^{-1}$ and an energy density of $2600 \text{ Wh}\cdot\text{kg}^{-1}$ [80,81]. Moreover, this technology may significantly reduce raw material costs and enhance operational safety compared with established Li ion chemistries. In addition, transition metal sulphide-based electrodes offer the advantage of high capacity via a conversion mechanism, in contrast to the intercalation mechanism of conventional cathodes [82]. An example is the primary Li-FeS_x battery, which features a lithium anode, an iron sulphide cathode, and lithium iodide in an organic solvent blend as the electrolyte.

The iron sulphide cathode can be in both the monosulphide (FeS) and disulphide (FeS₂) forms, following these reactions (discharging from left to right):

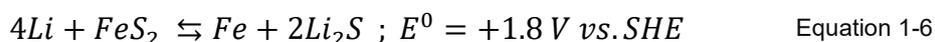
-Positive electrode:



-Negative electrode:



-Overall reaction:



Whereas sulphide-based solid electrolytes suffer from oxidative decomposition in contact with high voltage intercalation electrodes, the lower redox potential of FeS₂ possibly better agrees with the oxidative stability limit of thiophosphates. This constitutes an additional advantage for using FeS₂ as a cathode active material [83].

Jin *et al.* [81] synthesised carbon-coated iron disulphide (FeS₂@C) nanoparticles by the combined plasma evaporation method and sulphurisation. When employed in lithium-sulphur (Li-S) batteries, the Fe-S covalent interactions effectively suppress the formation and shuttling of polysulphides during charge–discharge cycles. Meanwhile, the carbon coating accommodates the volume expansion of the FeS₂ core and improves the electrical conductivity of this otherwise poorly conducting material. As a result, the dual-constrained FeS₂@C nanoparticles deliver high specific capacities and excellent cycling stability, even under ultrafast charging conditions ($10 \text{ A}\cdot\text{g}^{-1}$). Although the study primarily focused on Li–S systems, similar structural advantages could contribute to improved cycling stability in Li-ion batteries as well.

Dewald *et al.* [83] prepared FeS₂ nanoparticles with various size distributions and characterized their electrochemical performance in solid-state batteries. Reduced FeS₂

nanoparticle sizes showed a positive impact on the total capacity and rate capability, highlighting the potential of nanosized FeS₂ as cathode active conversion materials in solid state batteries.

Rechargeable magnesium batteries (RMB) are of particular interest due to the abundant magnesium resources and high volumetric capacity (3833 mAh·cm⁻³), as well as the uniform deposition of the anode. To achieve high-energy-density RMBs, researchers are focused on developing novel cathode materials that are not necessarily magnesium-containing [84]. This is because traditional Mg-based materials, often exhibit sluggish Mg²⁺ mobility and low operating voltages, which are major limitations for battery performance. Shen *et al.* [84] demonstrated a FeS₂ nanomaterial cathode with reversible capacity of 679 mAh·g⁻¹ at 50 mA·g⁻¹ and energy density of 714 Wh·kg⁻¹. This is superior to those of all previously reported metal chalcogenide cathodes in RMBs or hybrid batteries using a Mg metal anode. Notably, the as-assembled FeS₂–Mg battery can operate over 1000 cycles with a good capacity retention at 400 mA·g⁻¹.

Wu *et al.* [85] demonstrated a pulverisation phenomenon of ultrathin carbon-coated Fe_{1-x}S nanoplates during the first discharge process of a sodium-ion battery. This resulted in the formation of 5 nm Fe_{1-x}S nanoparticles, strongly embedded in the carbon matrix, which prevented aggregation. Additionally, researchers demonstrated both experimentally and through theoretical calculations that sodium adsorption and diffusion energy barriers were reduced. This sodium-ion battery, using iron sulphide as anode, showed a high reversible capacity (the amount of charge that can be repeatedly stored and released during cycling), reaching 610 mAh·g⁻¹ at 0.5 A·g⁻¹. A comparative summary of recent advancements in lithium–sulphur batteries using iron-based materials is provided in Table 1-3.

Table 1-3. Overview of research on lithium-sulphur batteries with iron-based components.

| Material | Specific capacity (mAh/g) | Current Density (A/g) | Coulombic Efficiency (%) | Energy density (W h /kg) | Ref |
|---|---------------------------|-----------------------|--------------------------|--------------------------|------|
| Iron disulphide core - nanoparticles with carbon shell | 862 | 0,5 | 100 | | [81] |
| Defect-rich iron disulphide nanoflowers on doped-carbon matrix | 841 | 0,89 | 94,5 | 500 | [82] |
| Iron disulphide nanomaterial | 679 | 0,05 | 97 | 714 | [84] |
| Ultrafine pulverized-Fe _{1-x} S nanoparticles with quantum size (≈5 nm) in a carbon matrix | 610 | 0,5 | 77,5 | - | [85] |
| Yolk-shell iron sulphide-carbon nanospheres | 545 | 0,1 | 70,2 | 438 | [86] |

1.2.5 Iron in Redox Flow Batteries

A redox flow battery (RFB) is an electrochemical device where at least one electrode comprises a solution of an electroactive material, and energy is generated/stored when the active species flow through the anode and cathode chambers separated by an ion-exchange (cation or anion) or microporous membrane and undergo electron-transfer reactions at inert electrodes [87]. As in the case of fuel cells, in a redox flow battery, the power and energy rating are decoupled.

Due to their low energy density, RFBs are not suitable for mobile applications. However, they offer several advantages, which can be summarised in four key features: moderate cost, modularity, transportability and flexible operation [88]. Their modular design also allows for potentially lower construction and maintenance costs compared to other energy storage technologies. This modular nature simplifies maintenance, as individual battery modules can be serviced separately. A major advantage is their operational flexibility during charge/discharge cycles: the batteries can be fully discharged without damaging the cells, an important advantage over established lead-acid technologies.

The use of iron in this kind of electrochemical device is motivated by its abundance, low-cost and non-toxicity. However, the electrochemistry is highly pH-sensitive, as the Hydrogen Evolution Reaction (HER) can occur, leading to the precipitation of Fe²⁺ ions

[89]. This remains one of the main challenges currently being addressed. Other relevant factors limiting the adoption of this technology include low current density, limited output efficiency (70-75%), and irreversible capacity loss due to electrolyte intermixing. A comprehensive description of how these challenges are being tackled, ongoing research trends, and future prospects of iron-chromium redox flow batteries can be found in ref. [90].

Despite the advantages this technology offers in terms of cost and abundance of its active materials, its commercialisation has not yet been achieved. In 2014, a US company called EnerVault installed a demonstration 250 kW / 1 MWh battery in California, funded by the US Department of Energy. According to the project's final report [91], the system showed an excellent degree of electrical and reactant flow uniformity. A key advantage highlighted by the DoE was that expanding energy storage capacity primarily involves enlarging tank sizes and increasing the volume of inexpensive electrolytes, making it an easily affordable endeavour. The additional cost for each extra kWh⁻¹ of energy storage can be less than \$100, making it a cost-effective option for long-duration applications. Despite these promising results, no further news related to EnerVault's projects or commercially available products has been found. In 2023, China Daily announced the installation and operation of the world's first megawatt-level iron-chromium flow battery in North China's Inner Mongolia, with a storage capacity of 6 MWh [92].

While several chemistries involving Fe exist, the most established is Iron-Chromium (ICRFB), which stores and releases electric energy through the following electrochemical reactions:

- Positive electrode:



- Negative electrode:



- Overall reaction:



The inexpensiveness and abundance of iron and chromium make them attractive choices for redox-active materials in ICRFB systems, contributing to an estimated cost of \$17 per kWh. This economic advantage provides a strong foundation for the development of cost-effective and competitive energy storage solutions [93]. The open-circuit voltage varies with the state of charge, primarily due to the dominance of different hydrated complexes of Cr³⁺ at different stages of operation.

Because chromium kinetics are inherently slow, catalysts such as bismuth or bismuth-lead on carbon have been developed to enhance performance of ICRFB. Che *et al.* [94] used a wet-chemistry method combined with self-polymerisation to deposit bismuth nanoparticles onto nitrogen-doped graphite felt. This modification not only enhanced electrochemical activity but also suppressed the hydrogen evolution reaction by forming intermediates through the reaction of Bi and H⁺ ions. Wan *et al.* [95] further addressed side reactions of HER by employing electrochemical purification of the electrolyte. This involved using a sacrificial electrode before cell operation to remove soluble contaminants, thereby reducing capacity fade rates by a factor of five.

Researchers at the Hong Kong University of Science and Technology investigated the impact of electrolyte flow design (serpentine vs interdigitated) on cell performance [96]. In a separate study, the electrochemical kinetics of the Fe²⁺/Fe³⁺ couple were compared to the V⁴⁺/V⁵⁺ couple using 1.0 M FeCl₂ and 1.0 M VOSO₄ solutions, respectively. The iron-based system exhibited smaller peak separation, indicating superior kinetics [97]. While oxidative pretreatment improves the electrochemical activity of carbon-based electrodes for the Fe²⁺/Fe³⁺ couple, the degree of graphitisation also plays a key role. In this study [98], less graphitized carbon felt showed higher charge transfer resistance, suggesting a trade-off between surface reactivity and conductivity.

Functionalisation of carbon felt can further improve the kinetics of the positive reaction in ICRFBs. Chen *et al.* [99], showed that SiO₂-decorated graphite felt, with its high density of surface oxygen-containing groups and hydroxyl functionalities, enhanced charge transfer during Fe²⁺/Fe³⁺ redox reactions. Additionally, treatment with silicic acid increased the effective surface area, boosting iron ion adsorption. These effects were corroborated by higher peak currents in cyclic voltammetry (CV) curves. However, the introduction of SiO₂ also resulted in larger redox peak potential separations, likely due to reduced graphitisation and electrical conductivity of the felt.

Xu *et al.* [100] reported a proof-of-concept of a novel redox flow battery consisting of a solid oxide electrochemical cell (SOEC) integrated with a redox-cycle unit. The charge/discharge characteristics were explicitly observed by operating between fuel cell and electrolysis modes of the SOEC along with 'in-battery' generation and storage of H₂ realized by an in situ closed-loop reversible steam-metal reaction in the redox-cycle unit. With Fe/FeO as the redox materials, the new storage battery can produce an energy capacity of 348 Wh per kg of Fe and round-trip energy efficiency of 91.5% over twenty stable charge/discharge cycles.

Gong *et al.* [101] described a new zinc-iron (Zn-Fe) RFB based on double-membrane triple-electrolyte design that is estimated to have under \$100 per kWh system capital cost is presented. This low cost is achieved through the use of affordable redox materials like zinc and iron, combined with high cell performance, reaching a power density of 676 mW·cm⁻². According to the authors, the engineering of the cell structure is essential for achieving this high-power density. The work, published in 2015, presents what might have been the cheapest RFB technology at that time, meeting the US DoE 2023 system capital cost target of \$150 per kWh. The standard potential is quite high compared to other iron-based RFBs, reaching 1.99 V. At a current density of 40 mA·cm⁻², the Zn-Fe RFB yields 74% energy efficiency.

The first alkaline redox flow battery (a-RFB) based on the coordination chemistry of cobalt with 1-[Bis(2-hydroxyethyl) amino]-2-propanol (mTEA) and iron with triethanolamine (TEA) in 5 M NaOH₂₄ is presented in this work [102]. Thanks to this chemistry, a cell voltage of 0.93 V in the charged state was obtained, while energy efficiencies reached 70-76% at current densities of 30 mA·cm⁻². The crossover of species through the Nafion cation exchange membrane during cycling was prevented due to the negative charge of the coordination compounds. A remarkable absence of crossover was observed for up to 30 charge-discharge cycles, surpassing commercial technologies based on vanadium. Moreover, no evolution of gases was detected.

New ferri/ferrocyanide - polysulphide (Fe/S) flow battery, which employs less corrosive, relatively environmentally benign neutral alkali metal ferri/ferrocyanide and alkali metal polysulphides as the active redox couples is presented by Wei *et al.* [103]. The catalyst on the polysulphide side consisted of cobalt nanoparticle -decorated graphite felt. Up to 99% coulombic efficiency and 74% energy efficiency was achieved, together with good cycling stability over numerous charge/discharge operations. An iron-cadmium redox flow battery (Fe/Cd RFB) with a premixed iron and cadmium solution is developed and tested by Zeng *et al.* [104]. In their study, researchers achieved a Fe/Cd RFB system with a coulombic efficiency and energy efficiency of 98.7% and 80.2% at 120 mA·cm⁻², respectively. They also reported good stability, with a capacity retention of 99.87% per cycle during the cycle test. The capital cost estimations were approximately \$108 kWh⁻¹ for 8-h energy storage.

To provide a concise comparison of various redox flow battery (RFB) systems discussed in this section, along with additional representative examples from the literature, Table 1-4 summarises their key electrochemical performance metrics. These include the redox chemistry, energy efficiency, current density, power density, and relevant references.

These parameters are commonly used to evaluate the performance of RFBs, as they directly reflect the cell's electrochemical behaviour and overall efficiency. Nevertheless, other factors such as durability (cycle life), long-term stability, operational safety, cost, and environmental impact are also important when assessing or selecting a battery technology. Although these aspects are not the focus of this table, they represent essential criteria for a comprehensive evaluation of RFB systems. This summary table highlights the variation in electrochemical performance and the current state of development across different RFB chemistries

Table 1-4. Summary of studies on iron redox flow batteries.

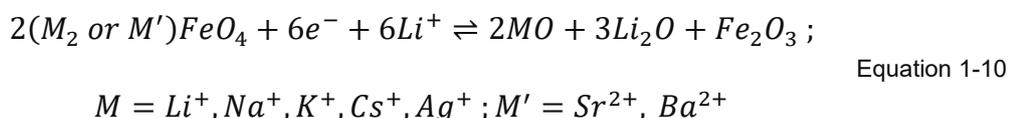
| Material | Energy Efficiency (%) | Current Density $\text{mA}\cdot\text{cm}^{-2}$ | Power Density $\text{mW}\cdot\text{cm}^{-2}$ | Ref |
|----------|-----------------------|---|---|-------|
| Fe-Cr | 80,7 | 320 | 665 | [93] |
| Fe-Cr | 80,5 | 480 | 1077 | [97] |
| Fe-Cr | 86,3 | 240 | 1077 | [97] |
| Fe-Cr | 84,8 | 320 | 1077 | [97] |
| Fe-Cr | 85,2 | 80 | 694 | [97] |
| Fe-Cr | 81,2 | 240 | 694 | [97] |
| Fe-Cr | 77,5 | 320 | 694 | [97] |
| Fe-Cr | 79,66 | 120 | | [99] |
| Fe-Zn | 76 | | 676 | [101] |
| Fe-Co | 71 | 30 | | [102] |
| Fe-S | 74 | 10 | | [103] |
| Fe-S | 70 | 30 | | [103] |
| Fe-S | 55 | 50 | | [103] |
| Fe-Cd | 80,2 | 120 | 125,4 | [104] |
| Fe | 50 | | 50 | [105] |
| Fe | 37 | 0,5 | | [106] |
| Fe | 41,5 | 0,25 | | [106] |
| Fe | 73 | 40 | 160 | [107] |
| Fe-Cr | 79,6 | 200 | 178 | [108] |
| Fe-Cr | 76,3 | 120 | 107 | [108] |
| Fe-V | 82 | 50 | | [109] |
| Fe-V | | 90 | 49 | [110] |
| Fe-V | 80,3 | 120 | | [111] |
| Fe-V | 75,2 | 160 | | [111] |
| Fe-V | 85,4 | 80 | | [111] |
| Fe-V | 89,4 | 40 | | [111] |

1.2.6 Super-Iron Batteries

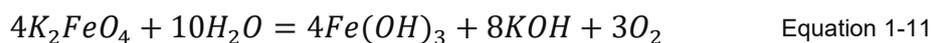
Hexavalent iron has been studied as a cathode material in both alkaline and Li-ion batteries. This 'super-iron' cathode, as first referred to by Licht *et al.* [112] due to its highly oxidized +6 valence state, offers several advantages. These include high intrinsic energy

supported by a 3-electron transfer process, environmental benignity, abundant raw materials (ferrates) and non-toxic discharge products.

Equation 1-10 below describes the overall reaction of a Li-ion super-iron battery



Although Licht's original article characterized a variety of salts, including Na_2FeO_4 , Li_2FeO_4 , Cs_2FeO_4 , $SrFeO_4$, Ag_2FeO_4 , $MgFeO_4$, $CaFeO_4$, $BaFeO_4$ and $ZnFeO_4$, it was K_2FeO_4 the one that garnered further interest in subsequent studies [113–117]. This is due to its high theoretical 3-electron charge capacity ($406 \text{ mAh}\cdot\text{g}^{-1}$) and superior stability in the solid state compared to the other salts. However, the redox potential of ferrate ion under acidic conditions exceeds the water stability voltage window, leading to the precipitation of ferric hydroxide and the decomposition of the cathode, as described in equation Equation 1-11 below.



In addition, the formation of Fe_2O_3 affects the reversibility of the charge transfer due to the poor electrical conductivity of the stable, passivating iron trivalent centres [118]. This limitation hampers the development of rechargeable devices in both aqueous and organic media. Consequently, research has focused on finding solutions to these major technical challenges, which hinder the development of commercial super-iron batteries.

Licht and coworkers studied the effect of adding transition metals such as manganese or silver oxide, to form composite with alkali or alkali earth super-irons [119–123]. The resulting composites enhanced charge/discharge transfer by improving charge transport mechanisms, while moving away from the current collectors the ferric-based passivated centres. The same group demonstrated the effectiveness of using of a platinum-based matrix as support to prevent passivation in non-aqueous media [112]. They also explored the deposition of a zirconia protective film for the same purpose in alkaline media [124,125].

This latter approach was further investigated by Zhang *et al.* [114], who used yttria-doped zirconia coatings as protective coatings for K_2FeO_4 electrodes in an alkaline electrolyte. The yttria doping significantly improved the electrochemical properties compared to the undoped zirconia-coated K_2FeO_4 cathodes, resulting in a more stable discharge plateau and up to 0.25 V higher discharge voltage. The coin cells, using zinc foil as anode,

achieved a discharge specific capacity of 169.8 mAh·g⁻¹ and a specific energy capacity of 201 Wh·kg⁻¹, outperforming cells with MnO₂-cathodes.

Farmand *et al.* [113] synthesised super-iron nanoparticles to enhance charge transfer using an optimised dry and cool ball-milling procedure that prevented thermal decomposition during milling. The nano-sized hexavalent iron cathodes demonstrated an 80% Coulombic efficiency at high discharge rates, compared to 53% for their micro-sized counterparts. Notably, the nanoparticles retained their hexavalent state up to 55% of capacity, while the microparticles transitioned to Fe₂O₃ at just 15% capacity.

Huang *et al.* [115] coated potassium ferrate with phthalocyanine via a coprecipitation process, demonstrating suppression of cathode degradation in an alkaline electrolyte, along with a higher open-circuit potential and enhanced charge transfer by inhibiting the formation of a trivalent iron layer. Wang *et al.* [116] investigated the effects of coating K₂FeO₄ with poly(3-hexylthiophene), resulting in a 13% increase in discharge capacity compared to the bare cathode, and retaining up to 22.6% more discharging capacity after 6 hours of immersion in an alkaline electrolyte.

Yan *et al.* [117] used small amounts of K₂SO₄ to dope K₂FeO₄, enhancing AAA-type alkaline super-iron batteries. The isomorphous substitution driven by SO₄²⁻ increased the discharge capacity by 10-30% compared to undoped cathode. Additionally, the cathode's stability was improved, showing a 10% higher shelf time.

In 2022, the Chinese automotive manufacturer GAC Group announced the launch of the first super-iron lithium battery, with 13.5% higher energy density than LiFePO₄-based devices, a 20% higher volume energy density and 10% higher capacity at -20°C [126].

1.2.7 Iron-Air Batteries

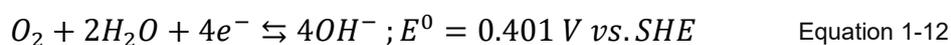
An iron–air battery is a type of metal–air battery that combines a metallic negative electrode with a low redox potential and an air positive electrode [127]. The theoretical voltage for this battery system, which is based on Fe(II)/Fe(0) followed by Fe(III)/Fe(II) redox reaction at the negative electrode and the oxygen reduction reaction (ORR) at the positive electrode, is around 1.28V. Therefore, the theoretical specific energy is around 1000 Wh·kg⁻¹.

The electrode reaction of iron in an alkaline electrolyte occurs entirely in the solid state without involving a deposition–dissolution process due to the low solubility of the discharge products, which are iron hydroxides and iron oxides. This property of iron is a disadvantage for a primary battery but a potential advantage for a secondary battery, as it eliminates concerns about the shape change phenomenon of the negative electrode

over cycles. Consequently, this battery system is considered an advanced secondary battery.

The cell reactions are:

-Positive electrode:

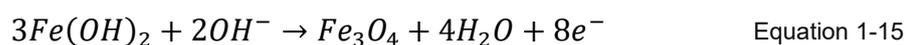
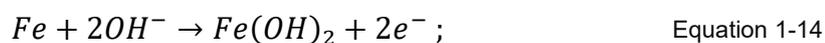


-Negative electrode:

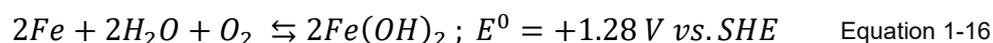
Charging



Discharging



-Overall reaction:



Unlike zinc in the zinc–air battery, iron is less prone to forming dendrites with repeated charge–discharge cycles from aqueous electrolytes, which can cause severe changes in the shape of the electrodes during prolonged cycling. Iron is also a very low cost, environmentally friendly resource (recycling being a well-established process), although the strong alkaline electrolyte requires adequate protective equipment and safety procedures. The advantages of the iron–air battery include its moderate energy density of 50–75 Wh·kg⁻¹ and a cost below US\$100(kWh)⁻¹ [128].

One of the major limitations of this technology is the significant extent of hydrogen evolution reaction that may occur during charging. In alkaline media, hydrogen is generated by the electrochemical reduction of water ($2H_2O + 2e^- \rightarrow H_2 + 2OH^-$), and its evolution potential is close to the Fe²⁺/Fe redox potential [127]. Thus, suppressing hydrogen evolution is key for optimal charge–discharge efficiency. In addition, the undesirable hydrogen evolution during the storage of the battery cell at the charged state converts iron to iron hydroxide and brings the self-discharge of the battery.

Presently, almost all commercial metal–air batteries use zinc as negative electrode due to its higher capacity and working voltage compared to iron. However, a spinoff from MIT called Form Energy, co-founded by Mateo Jaramillo, ex-Tesla’s Vice President of

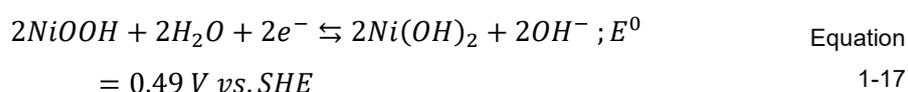
Stationary Energy Products and Programs, claims to have developed a commercial iron-air battery optimized to store electricity for 100 hours at a competitive cost [129]. ArcelorMittal invested a total of \$42.5M in the company through its innovation fund XCarb®.

Additionally, at the end of 2023, it was announced that Xcel Energy, a major U.S. electric utility company, would receive a \$70M grant from the U.S. Department of Energy to partially fund a demonstration-scale 10MW / 1000 MWh system based on iron-air batteries at two retiring coal plant sites [130].

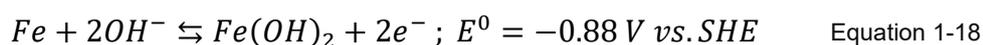
1.2.8 Nickel-Iron Batteries

Nickel-Iron aqueous alkaline batteries are probably the most investigated aqueous energy storage system using iron as anode material [131]. Early Ni-Fe batteries, which were patented and made commercially available by Thomas Edison at the beginning of the twentieth century, have been used for more than 70 years in both stationary and mobile energy storage applications in Europe and the United States. The reactions taking place in a Nickel-Iron battery are as it follows (discharging from left to right):

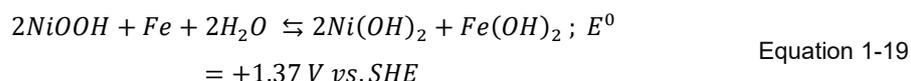
-Positive electrode:



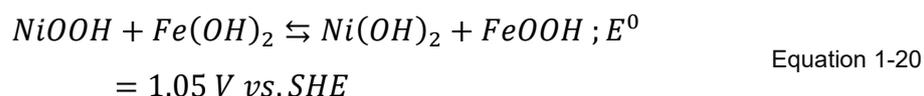
-Negative electrode:



-Overall reaction:



Under deep discharge, Ni-Fe will undergo a further discharge reaction at a lower voltage than that of the reaction described above, namely:



One of the main advantages of these batteries is the abundance of both iron (anode) and water (electrolyte), which makes the system a lot cheaper and safer compared to non-aqueous chemistries. Due to the high reversibility under alkaline conditions of the previously mentioned reactions, these batteries have typically long charge-discharge cycle life.

Ni-Fe batteries work well at low or moderate rates of discharge, but they do not perform well for high discharge rates applications [132]. This is because high discharge rates enhance the passivation of electrodes through the formation of $Fe(OH)_2$ (or $FeOOH$ under deep discharge), an insulating layer, at the surface of the electrode. This prevents the active material from being fully utilized. It has long been known that the presence of sulphides in the iron electrode can help mitigate the formation of a passive layer. Manohar *et al.* [133] studied how the addition to the iron electrode of two sulphide additives (sodium sulphide and iron sulphide) affected the passivation. With the addition of sodium sulphide, the capacity at C/1 rate was 8 times more than the same electrode without that additive. Moreover, by incorporating iron sulphide at 5w.% the capacity was increased by 18 times compared to the reference without additives. According to the researchers, the experiments at 3C rate, with an electrode utilisation of almost $0.2 \text{ Ah}\cdot\text{g}^{-1}$ demonstrated the highest discharge rate ever reported with iron electrodes.

The open-circuit potential of a charged iron electrode under alkaline conditions ($E^0 = -0.88 \text{ V vs. SHE}$) is more negative than hydrogen evolution potential under the same pH conditions ($E^0 = -0.83 \text{ V vs. SHE}$). This causes the corrosion of the iron electrode [134], low-charging efficiency (less than 60% in commercial alkaline nickel-iron batteries [133]), self-discharge of the battery and the loss of water from the electrolyte. Furthermore, reduction of hydrogen evolution is a must in order to enable commercial Ni-Iron batteries that are sealed and that don't require special maintenance requirements. Different strategies have been considered to overcome this problem. Posada *et al.* [135] studied how the addition of lithium hydroxide and potassium sulphide to the electrolyte impacts the performance of a Ni-Fe cell. They concluded that the addition of potassium sulphide mitigates efficiently the evolution of hydrogen on the iron electrode. On the other hand, adding additives to the iron electrode like bismuth, sulphur [136–139] and copper-containing [140] species can also be effective in reducing the hydrogen evolution reaction by increasing the activation barrier for water decomposition.

Energy and power densities of Nickel-Iron batteries are rather poor ($\sim 30\text{-}50 \text{ Wh}\cdot\text{kg}^{-1}$ and $\sim 3\text{-}50 \text{ W}\cdot\text{kg}^{-1}$) compared to Li-Ion technology, limiting their commercial use to niche applications like industrial forklift trucks, open-pit mining vehicles or railroad signalling [132,141]. In the past, Ni-Fe batteries were the focus of research by multiple industry players such as Eagle-Picher Industries, SAFT, VARTA or SAB-NIFE. Nowadays, researchers are exploring new chemistries with potentially higher commercial interest, such as Ni-metal hydrides, Ni-Cd and Li-Ion. However, Ni-Fe technology is still under consideration by research groups working on radical new battery concepts like

3D-printed compressible quasi-solid batteries [142] or nickel-iron 'battolyser', i.e. a combined battery and electrolyser [143].

1.2.9 Other Uses

While this review has primarily focused on the electrochemical properties of iron and its derived electrochemical devices, steels, which are a broad family of iron alloys, are extensively used as structural materials in many energy storage applications. Manufacturers such as ArcelorMittal offer a wide range of steel grades specifically developed for use in electric vehicle battery packs [144]. The increasing demand for steel in this area is driving the development of new commercial grades with improved mechanical properties. In addition, new coating solutions, particularly those based on nickel, are being developed to ensure effective electrical contact between the steel casing and the aluminium or copper tabs [145,146]. Graphene coatings are also being explored as a protective layer to prevent corrosion of the nickel coatings [147,148].

Another active area of research focuses on advanced laser welding techniques. These are aimed at creating robust joining technologies that meet industrial standards, which is key for maintaining the quality of electrical interconnections between the current collector and the steel battery can [149–152].

Figure 1-6 illustrates the different joining operations required in battery-pack manufacturing. The figure shows typical welding locations for cylindrical, pouch and prismatic cells, as well as examples of interconnections between cells and current-collecting busbars. These schematics highlight the variety of interfaces that must be electrically and mechanically joined within a battery module. They emphasise the demanding role played by metallic materials, including steels, in the structural and electrical integration of battery systems. Each joining interface must simultaneously provide low electrical resistance, high mechanical integrity and long-term corrosion resistance. From a physical standpoint, this places strict requirements on the thermal, mechanical and surface properties of the metals involved.

Thanks to its abundance, mechanical strength, and corrosion resistance, steel is considered a strong candidate material for components used in hydrogen transport and storage systems. The European Hydrogen Backbone project, which aims to reach 60,000 km of hydrogen pipelines by 2040 [153], considers steel as the primary material for high-pressure transmission lines. Experts suggest that hydrogen pipelines will be quite similar in design to those used for natural gas [154]. However, transporting hydrogen at high pressure introduces challenges, particularly due to hydrogen embrittlement, a phenomenon that can weaken steel when hydrogen atoms diffuse into

the material [155–157]. This has increased interest in the development of leakage monitoring technologies, with particular attention given to fibre optics and nanosensors [158–162].

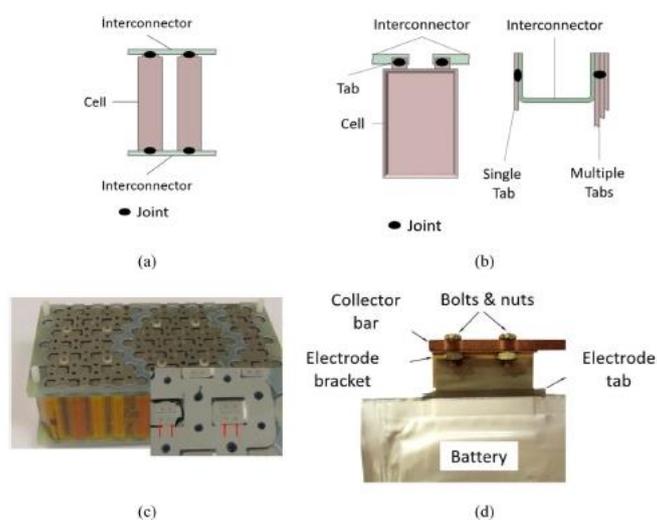


Figure 1-6. Overview of joining tasks in battery applications: schematic depiction of the joining location (a) for cylindrical cells and (b) for pouch or prismatic cells; (c) battery module consisting of cylindrical cells directly connected by a large busbar, reproduced from [163]; (d) single pouch cell interconnected to a copper collector bar, reproduced from [164] with permission from Elsevier. The figure is reproduced from [34] under the CC BY-NC-ND license.

At the same time, new surface treatments and coatings are being investigated to reduce hydrogen diffusion into steel. These include metallic coatings such as vanadium, titanium, and zinc-nickel, as well as compound coatings made from nitrides, oxides, and carbides [165–167]. A particularly notable study by Shi *et al.* [168] showed that graphene-coated pipe steel exhibited a nearly 50-fold reduction in hydrogen permeability, providing outstanding resistance to hydrogen embrittlement.

The key contributing factors to hydrogen embrittlement include high stress concentrations in the steel, the presence of diffusible hydrogen, and the chemical composition of the material. This problem is especially pronounced in high-strength steels, as highlighted in several studies [169–171], although the underlying mechanisms are not yet fully understood.

From a chemical perspective, it is known that certain elements such as carbon (C), silicon (Si), phosphorus (P), and sulphur (S) can increase susceptibility to hydrogen embrittlement. In contrast, the presence of nickel (Ni), aluminium (Al), and molybdenum (Mo) tends to have a beneficial effect [172–175]. For instance, Kwon *et al.* [173] demonstrated that in Mn-B steels, an increase in carbon content leads to higher embrittlement up to a certain limit. Similarly, Koyama *et al.* [176] observed that this effect occurs for carbon levels below 0.3 wt.%, but higher concentrations do not lead to further

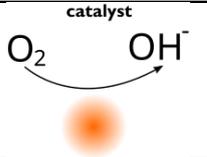
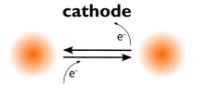
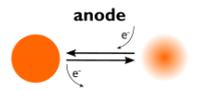
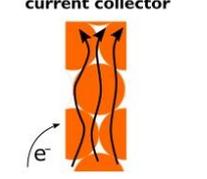
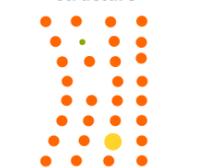
increases in susceptibility. Nickel is particularly important in sour-resistant line pipe steels and may also contribute to the development of durable steel structures for hydrogen storage. Mohrbacher *et al.* [172] reported that nickel can limit austenite grain growth, which in turn reduces the amount of diffusible hydrogen and enhances resistance to embrittlement.

Stainless steels, which contain high levels of chromium to prevent corrosion, form a distinct group within the steel family and offer unique advantages for electrochemical applications, particularly as current collectors. While current collectors in lithium-ion batteries are typically made of aluminium or copper due to their corrosion resistance and their ability to be processed into thin, electrochemically stable foils. However, stainless steel presents several advantages over these conventional materials. It can function as a current collector for both positive and negative electrodes, has a lower specific density than copper, and, under specific conditions such as in nickel-cobalt-aluminium (NCA) Li-ion batteries, can offer higher corrosion resistance than aluminium [177]. Stainless steel can also be plated with a thin nickel layer to reduce contact resistance and further improve corrosion resistance [178,179].

In fuel cell technology, stainless steel bipolar plates are already commercially available and are offered as lower-cost alternatives to graphite plates by companies such as Outokumpu (Finland) [180] and Waelzholz (Germany) [181]. Stainless steel plates are more durable than graphite, making them better suited to handle the dynamic conditions within operating fuel cells [182]. Additionally, the higher ductility of steel simplifies the formation of flow channels during manufacturing, which helps reduce production costs, although a corrosion prevention method needs to be applied to the surface.

To conclude, iron plays a multifaceted role in modern energy storage systems due to its abundance, cost-effectiveness, and favourable electrochemical properties. Table 1-5 summarises the various roles of iron across different technologies, highlighting key aspects such as the specific form of iron involved, the technical challenges associated with each application, the level of technological maturity, and relevant references.

Table 1-5. Overview of iron's roles in various energy storage technologies, including the specific form used, associated technical challenges, technological maturity, and key references.

| Role | Technology | Iron form | Technical Challenges | Technological Maturity | Refs |
|--|--------------------|--|---|--------------------------------------|-----------|
|  catalyst | PEM Fuel Cells | Porous, Fe ^{2+/3+} -containing compound | Demetallation | Research status | (52-68) |
|  cathode | Li-Ion | LiFePO ₄ | Moderate electrical conductivity | Massively adopted by the market | (69-78) |
| | Li-Sulphur | FeS ₂ | Low electric and ionic conductivities Volumetric changes Undesired discharge products | Research status | (79-85) |
| | Redox-Flow | Fe ²⁺ -based salt or complex | Cross-contamination Hydrogen Evolution Reaction Slow kinetics | First MWh plant operative since 2023 | (86-111) |
| | Super Fe | Fe ⁶⁺ -based salt | Irreversible reduction | First commercial battery announced | (112-126) |
|  anode | Redox-Flow | Fe | Cross-contamination Hydrogen Evolution Reaction Slow kinetics Low cell voltage | Research status | (86-111) |
| | Iron-air | Fe | Hydrogen Evolution Reaction | Pilot demonstration | (127-130) |
| | Ni-Fe | Fe | Hydrogen Evolution Reaction | Obsolete | (131-143) |
|  current collector | PEM Fuel Cells | Stainless steel foil | High mass density | Commercially available | (174-176) |
| | Li-Ion | Stainless steel foil, bare or Ni-plated | High mass density | Commercially available | |
|  structure | Li-Ion | Steel battery pack | High mass density | Commercially available | (143-156) |
| | Hydrogen transport | Steel pipe | High mass density Hydrogen embrittlement | Commercially available | (177-198) |
| | Hydrogen storage | Steel vessel/ pipe | | | (199-209) |

1.2.10 Research Gap

As outlined in the first part of the literature review, iron-based materials have been widely studied for their applications in electrochemical devices. However, the potential of Kish graphite, a byproduct of the steel industry, remains underexplored. Kish graphite contains intercalated iron compounds that may act as catalysts for the ORR, which is the rate-limiting step in proton exchange membrane fuel cells. Despite this potential, its catalytic performance and broader electrochemical applications have not been systematically studied.

Unlike deliberately synthesized iron–carbon composites, Kish graphite naturally contains iron-rich inclusions and intercalated iron compounds embedded within a highly graphitized carbon matrix. This combination may provide catalytically active iron sites while simultaneously offering high electrical conductivity and structural stability, both of which are advantageous for ORR catalysis. While other iron-containing carbon materials have been shown to exhibit ORR activity, Kish graphite is distinguished by its abundance, low cost, and minimal processing requirements, as well as by the intimate coupling between iron species and a conductive graphitic host. Despite these potentially favourable attributes, its catalytic performance and broader electrochemical applications have not yet been systematically evaluated or quantitatively benchmarked against established iron-based ORR catalysts.

In addition, Kish graphite can be used as a raw material for producing reduced graphene oxide (rGO), a material with well-established applications as a catalyst support or as an electrode in supercapacitors.

The experimental and machine learning methodologies required for these investigations are presented in Chapter 2, including the synthesis and processing of Kish graphite-derived materials, the electrochemical characterisation techniques, and the machine learning framework.

To address this gap, Chapter 3 explores the use of expanded Kish graphite as a catalyst for the ORR. To enhance the oxygen accessibility to iron sites, Kish graphite is expanded following the method described in Section 2.2.3. Furthermore, rGO is synthesized from Kish graphite using the method outlined in Section 2.2.2. Platinum is deposited onto rGO, and its performance is compared to a commercial benchmark catalyst, 40% Pt on Vulcan XC72 (Sigma-Aldrich). Electrochemical characterisation is complemented by capacitance measurements of rGO samples reduced under different temperatures and durations.

Given the overall motivation of this thesis to apply machine learning in materials research, Chapter 3 also includes an analysis of the experimental noise associated with these types of measurements at different replication levels. This experimentally quantified noise will serve as the foundation for the implementation of physics-informed Gaussian Process models in Chapter 4, whose objective is to determine the optimal reduction conditions for maximising rGO capacitance. The motivation for this study is introduced in Section 1.3 below, which presents a literature review of graphene supercapacitors. This review also provides the basis for the work described in Chapter

4, which focuses on developing a machine learning model for graphene capacitance with an emphasis on physical correctness and adaptation to data-scarce conditions.

1.3 Literature Review: The Role of Graphene in Supercapacitors

Graphene is a 2D material consisting of carbon atoms forming a honeycomb structure based on sp^2 bonds. Since the experimental realisation of graphene stability at room temperature and its outstanding electrical properties in 2004 [44], research on graphene has been at the forefront in universities and research centres all over the world. In fact, the first publication about graphene from the Nobel-awarded scientists André Geim and Konstantin Novoselov [44] is one of the most highly cited articles of all time. The ambipolar field effect and the flow of graphene electrons as if they had no mass have motivated the development of multiple applications for graphene in the field of electronics and optoelectronics [183,184]. Moreover, a cascade of high-impact papers demonstrated that graphene is the material with the highest thermal conductivity [185], highest tensile strength and Young Modulus [186] and impermeable to all liquids and gases [187,188]. However, graphene is not impermeable to protons (H^+) and other cations, a property that has motivated research into its use as an ion-selective membrane for fuel cells [189].

With the goal of translating the outstanding properties of graphene into real innovative commercial products and solutions, in 2013 the European Commission launched the Graphene Flagship, a €1 billion research initiative to support both graphene research and the industrialisation of a graphene-enabled market [190–194]. One of the most important applications highlighted by the Graphene Flagship for graphene is in the field of energy storage, and specifically batteries and supercapacitors. Figure 1-7 illustrates the relationship between the cumulative average growth rate (CAGR) and the current market revenue for graphene materials across different application sectors. The y-axis shows the CAGR, while the x-axis presents the revenue in \$US millions on a logarithmic scale. Among the surveyed markets, the aerospace sector exhibits the highest revenue, whereas the battery sector shows the highest projected growth.

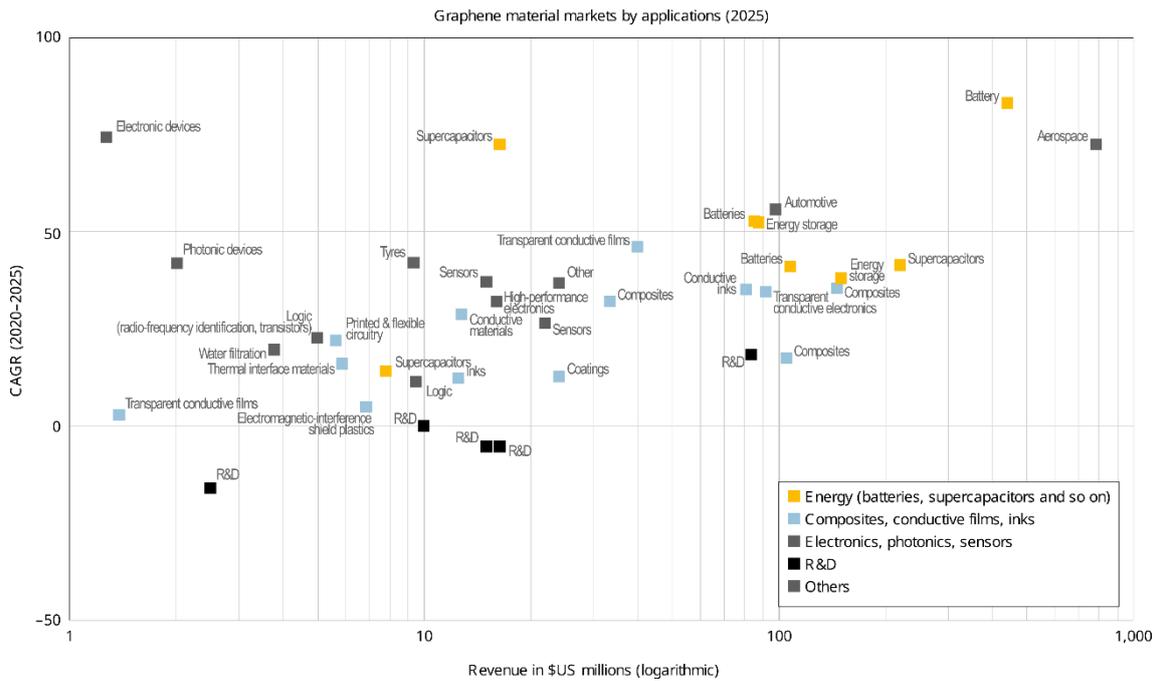


Figure 1-7. Graphene markets by application by 2025, reproduced with permission from Reiss *et al.* [194].

As a two-dimensional material that is conductive and carbon-based, graphene is quite interesting in supercapacitors (SCs). In these devices, the energy is stored on the electrostatic double layer at the interface between the electrode and the electrolyte. Contrarily to batteries, no redox reactions take place, providing a system that can be charged and discharged much more quickly than batteries, which enables devices with much higher power density. In addition, there is no risk of overcharging and fire hazard is negligible.

Supercapacitors show high lifetime, keeping a good performance after hundreds of thousands of charging/discharging cycles and with low maintenance costs. However, as the charge is limited by the electrode surface that is available, they have much lower energy density compared to batteries. They have also higher tendency to self-discharge [195].

Commercial SCs can be found in applications where there is a need for fast delivery of power such as in industrial machinery, urban transportation, power tools, communication devices, robotics and motor start up systems [192]. They also can have a key role in grid electricity networks for applications such as rectifying sudden frequency and voltage changes across the network or adjusting unanticipated asymmetries between generation and load [196].

The recent blackout affecting Spain and Portugal in April 2025 was one of the most severe in recent European history, causing widespread disruption to transportation,

communications, and essential services [197,198]. The incident, triggered by the sudden loss of 2.2 GW of generation in southern Spain, causing an abrupt imbalance between supply and demand, highlighted the pressing need for rapid-response energy storage solutions, such as supercapacitors (SCs), which could have injected power instantly to stabilize grid frequency and voltage during the disturbance. SCs offer unique advantages in addressing such challenges, particularly in four key applications related to grid stability and resilience [196]:

- Primary response: to guarantee supply under continuous and sudden frequency and voltage changes across the network
- Secondary response: to compensate anticipated and unexpected imbalances between load and generation
- Black start: when restoring power plant operations after network outage without external power supply
- Power quality: to protect on-site load against short-duration power loss or variations in voltage or frequency

The market volume for these devices has surpassed \$1 billion in the recent years [199], with suppliers like Maxwell, Panasonic, Skeleton, Ioxus and JM Energy. Among these, Skeleton outperforms the cells from the other suppliers according to a study from the Naval Surface Warfare Centre [200], with an Equivalent Series Resistance (ESR) of less than half of any of the other cells and a power density more than 10-kW/kg higher than the others. In fact, Skeleton is, out of these 5 suppliers, the only one that claims the use of graphene as electrode material.

Specifically, Skeleton describes a patented formulation of a curved graphene [201–204] with superior properties with respect to activated carbons [205], delivering 60% higher energy storage capacity and double power performance [206]. The company has its production facility in Grossröhrsdorf (Germany), which is the largest supercapacitor manufacturing capacity in Europe. Moreover, a new facility is being built in Markranstädt, after a strategic partnership with Siemens [207]. In total, the company has raised over €200 M investment to develop and commercialize graphene-based energy storage technologies.

The main advantages of graphene materials with respect to activated carbons are a higher usable surface area for the electrolyte, due to a more favourable mesoporous structure [208,209] arising from the arrangements of graphene layers around each other, and, for rGOs, also from their holey-structure, as well as higher electrical conductivity [210] and greater mechanical and thermal stability [191]. Commercially available

supercapacitors based on activated carbon reach energy density in the range of 7.8-20 Wh·L⁻¹ [200,211], while graphene-based flexible supercapacitors have already demonstrated 88 Wh·L⁻¹ [212] in the literature. In terms of material costs, activated carbons are under 15\$·kg⁻¹ while graphene is more than ten times more expensive [192], which is one of the main barriers for its commercial development.

With respect to graphene research for supercapacitors, the principal aspects to consider regarding its physical attributes that contribute to a good performance as electrode are the following [213,214]:

- Large active area. The capacitance of a material is given by $C = \frac{\epsilon_0 \epsilon_r A}{d}$ where A is the electrode surface area, ϵ_0 and ϵ_r are the vacuum permittivity and material absolute permittivity constants, respectively and d the distance between the electrolyte's ions and the electrode's walls. Thus, increasing A is a good strategy to maximize the capacitance of the electrode material. This parameter is typically measured by flowing N₂ gas into the porous material and characterizing the adsorption/desorption isotherms, using the Brunauer-Emmett-Teller (BET) to estimate the Specific Surface Area of the material.
- Good electrical conductivity, to minimize the loss of power from internal resistance. This is often estimated indirectly by calculating the I_D/I_G ratio from the Raman spectra of the graphene material, as an indicator of the density of defects in the crystalline structure. In addition, electrochemical performance features can be used as an indirect indicator of the higher or lower electrical conductivity from the measured electrode materials.
- Optimal pore size distribution. Micropore volume enhances the total capacity while hierarchical pores in the micro- and mesoporous region facilitate rapid ion transport [214]. This is also extracted from the N₂ adsorption/desorption isotherms, combining quenched solid density functional theory for micro- and meso- pores <50 nm and Barrett-Joyner-Halenda (BJH) method for macropores >50 nm.
- Surface wettability, to enhance pore flooding. This can be estimated with contact angle measurements using the sessile-drop method.

A detailed explanation about the experimental techniques and methods used to characterize graphene physical properties and electrochemical performance is described in Chapter 2.

Table 1-6 below summarises results from recent works of the capacitive performance of graphene-based electrodes. As it can be extracted from the table graphene, in its

multiple forms, outperforms the specific capacitance of the commercial activated carbon reference.

Table 1-6. Summary of specific capacitance of different graphene materials (partially extracted from [214–216])

| Sample | Electrolyte | Current Density | Specific Capacitance (F/g) | Ref. |
|---|--------------------------------------|-----------------------|----------------------------|-------|
| RGO | 1 M H ₂ SO ₄ | 1 mV s ⁻¹ | 198 | [217] |
| N- doped holey graphene | 6 M KOH | 0.5 A g ⁻¹ | 250 | [218] |
| CF/graphene foam | 3 M KOH | 0.5 A g ⁻¹ | 162 | [219] |
| 3D microporous graphene | 1M Na ₂ SO ₄ | 1 A g ⁻¹ | 202 | [220] |
| Graphene nanoballs | 1 M H ₂ SO ₄ | 5 mV s ⁻¹ | 206 | [221] |
| N- doped graphene | 6 M KOH | 0.5 A g ⁻¹ | 255 | [222] |
| 3D N and B co-doped graphene | 1 M H ₂ SO ₄ | 1 A g ⁻¹ | 239 | [223] |
| RGO | 6 M KOH | 0.5 A g ⁻¹ | 436 | [224] |
| N- doped graphene aerogels | 1 M H ₂ SO ₄ | 0.2 A g ⁻¹ | 223 | [225] |
| N-doped graphene | 0.5 M H ₂ SO ₄ | 0.2 A g ⁻¹ | 286 | [226] |
| N-doped graphene | 6 M KOH | 0.1 A g ⁻¹ | 301 | [227] |
| 3D porous RGO | 1 M H ₂ SO ₄ | 1 A g ⁻¹ | 284.2 | [228] |
| Holey graphene | 1 M H ₂ SO ₄ | 0.5 A g ⁻¹ | 250 | [229] |
| RGO hierarchical structure | 6 M KOH | 1 A g ⁻¹ | 429.7 | [215] |
| Hole graphene nanosheets | 6 M KOH | 50 A g ⁻¹ | 170 | [230] |
| Nitrogen-doped holey graphene | 2 M H ₂ SO ₄ | 0.3 A g ⁻¹ | 343 | [231] |
| N and F co-doped holey graphene hydrogel | 6 M KOH | 1 A g ⁻¹ | 345.4 | [232] |
| P and N doped holey graphene | 1 M H ₂ SO ₄ | 0.5 A g ⁻¹ | 235.5 | [233] |
| N-doped porous carbon/holey graphene | 6 M KOH | 1 A g ⁻¹ | 323.5 | [234] |
| Holey graphene / Ppy films | 6 M KOH | 1 A g ⁻¹ | 438 | [235] |
| Holey graphene / PANI / graphene films | 1 M H ₂ SO ₄ | 1 A g ⁻¹ | 804 | [236] |
| Self-assembled graphene hydrogel | 5 M KOH | 1 A g ⁻¹ | 160 | [237] |
| Chemically modified graphene | 5.5 M KOH | 1.3 A g ⁻¹ | 135 | [210] |
| Graphene aerogel mesoporous carbon | 1 M H ₂ SO ₄ | 1 mV s ⁻¹ | 226 | [238] |
| Holey graphene framework | 6 M KOH | 1 A g ⁻¹ | 310 | [239] |
| Reduced holey graphene oxide hydrogel | 1 M H ₂ SO ₄ | 1 A g ⁻¹ | 283 | [240] |

| Sample | Electrolyte | Current Density | Specific Capacitance (F/g) | Ref. |
|--|------------------------------------|-----------------------|----------------------------|-------|
| Thermally expanded graphene oxide | 5.5 M KOH | 0.1 A g ⁻¹ | 264 | [241] |
| Non-stacked reduced graphene oxide | 6 M KOH | 1 A g ⁻¹ | 237 | [242] |
| High density porous graphene | 6 M KOH | 0.1 A g ⁻¹ | 238 | [243] |
| Ultrahigh-level oxygen functionalized graphene | 1 M KOH | 1 A g ⁻¹ | 330 | [214] |
| Commercial Activated Carbon Norit™ | 1 M H ₂ SO ₄ | 0.5 A g ⁻¹ | 91 | [244] |
| Commercial Activated Carbon Norit™ | 6 M KOH | 0.5 A g ⁻¹ | 86 | [244] |

Beyond supercapacitors, there are dozens of articles related to the use of graphene and its derivatives as electrodes for a large variety of electrochemical devices, including Li-ion [245,246], Li-S [247], Na-Ion [248] and redox flow batteries [249].

1.3.1 Machine Learning for the Design of Experiments

Design of Experiments (DoE) consists in an ensemble of methods supported by statistics with the aim of guiding the execution of experiments to detect relevant variables with the minimum number of iterations possible. The use of DoE techniques can result in cheaper and easier development cycles in research and product innovation [250]. An adequate selection of DoE technique facilitates the process of inferring patterns between the initial set of parameters and the outputs from the experiment. Literature on DoE dates back to 1936, with the application of statistical methods for planning experiments in agriculture research [251]. Since then, several designs have been consolidated, namely: Response Surface Methodology (RSM) [252], Robust Parameter Design (RPD) [253] or Latin Hypercube Sampling (LHS) [254].

Sequential Experimentation (SE) is a concept that can be applied to all types of DoEs and whose principle is to perform a series of experiments in a sequential manner, using the results of the previous experiments to guide the following iterations.

Bayesian Optimisation (BO) is an increasingly widely used statistical tool that uses the principles of Design of Experiments and Sequential Experimentation with the support of Machine Learning methods. The goal is to minimize the number of experiments required to get materials with a set of given desired properties, using Bayesian inference to drive the next steps. In addition, it also serves to gain knowledge about the relationships

between the material properties and its synthesis parameters in a much more efficient way than following traditional methods which are not supported by Artificial Intelligence.

BO is, in essence, a probabilistic framework for the global optimisation of a 'black-box' function (i.e., an objective function). This is achieved by sequentially selecting inputs and observing the corresponding outputs, without requiring explicit knowledge of the function's analytical form. A 'black-box' function refers to a system where outputs can be obtained from inputs, but the exact form of the function is unknown. In materials science, such functions are common, as the relationship between synthesis parameters and resulting physicochemical properties is often poorly understood. Thus, the discovery of new materials is limited by the 'lack of light' in the chemical design space. This issue is particularly pronounced in the development of nanomaterials and soft materials, where the atomic or molecular spatial arrangements exhibit greater degrees of freedom than in bulk, crystalline materials. As a result, small variations in synthesis conditions can lead to significant changes in material properties and reduced reproducibility.

Instead of evaluating all possibilities in a predefined synthesis parameter space following traditional grid search and subsequent resource-intensive experiments/simulations, Bayesian Optimisation enables the identification of the material candidate by iteratively updating a surrogate model to select the next experimental conditions to try [255]. This provides a framework for the Design of Experiments based on machine-learning that minimizes the number of experimental tests and/or simulations needed compared to traditional grid approaches.

In this paper [256], the researchers apply BO to find the chemical compound with the lowest thermal conductivity for thermoelectric applications. Firstly, a prior model was built by taking an initial dataset of five compounds and its thermal conductivities determined by first-principles methods. Two predictors (components of \mathbf{x} vector) were selected: crystalline volume per atom, V and density, ρ . An advantage of choosing these is that both physical descriptors are available in most of the experimental or computational databases like ICSD [257], Atomwork [258] or Materials Project [259]. Gaussian Process regression was used with low Lattice Thermal Conductivity (LTC) as objective function. As the acquisition function, Probability of Improvement was selected.

The results from the iterative process were compared with a traditional random search method for a total of 50 iterations, and the average number of observed compounds required for finding the one with the lowest thermal conductivity from a population of 101 was evaluated. Using random search, the average number was 55 while the number was only 11 for the case of BO. However, the scientists demonstrated that using V and ρ as

predictors was not an optimal choice: just by extracting the first and second lowest thermal conductivity compounds out of the 101 population, the Bayesian optimisation procedure performed worse than random search (65 compounds needed to find the one with lowest conductivity versus 50, respectively).

The reason, they argued, was that the new objective compound was an outlier when the thermal conductivity is modelled only with V and ρ . By adding extra predictors that are related to the chemical nature of the compounds the researchers were able to overcome that issue, reducing the average number of compounds evaluated to just 19. Once this tuning of the Bayesian Optimisation framework was complete, they applied the same methodology to explore compounds with low conductivity from a population of 54,779, finding 221 materials with very low lattice thermal conductivity. Among them, two (K_2CdPb and $Cs_2[PdCl_4]I_2$) presented a narrow electronic band gap, which makes them specially interesting for thermoelectric applications.

Machine Learning can also be applied to combine data from first-principles calculations and high-throughput experimentation in a closed-loop optimisation framework, as demonstrated by Sun *et al.* [260]. In this research, the authors explored the quasi-ternary A-site compositional space of $Cs_xMA_yFA_{1-x-y}PbI_3$, defined by the Cs-MA-FA compositional simplex and discretised at 1% resolution, resulting in more than 5,000 possible compositions. Using a physics-informed Bayesian optimisation approach that incorporates density functional theory-derived phase stability as a probabilistic constraint, they were able to identify highly stable multi-cation lead iodide perovskites by experimentally sampling only 1.8% of the total compositional space.

A physics-informed Machine Learning strategy was adopted, in which Density Functional Theory-derived Gibbs free energies of mixing were incorporated as probabilistic constraints instead of applying a hard constraint boundary. This ensured that the search was guided toward thermodynamically and structurally stable compositions rather than relying solely on experimental performance data.

The sparse set of 91 DFT ΔG_{mix} values is first interpolated across the full ternary Cs-MA-FA (Cs-methylammonium-formamidinium) space using a Gaussian Process surrogate model, whose predictive variance reflects both the limited coverage of the DFT dataset and the intrinsic 0.025 eV per formula-unit precision of the calculations arising from configurational variability. The resulting mean $\Delta G_{mix}(Q)$ is then transformed into a probability of phase stability via a logistic function, $P(\Delta G_{mix}(Q); \beta_{DFT}) = 1/(1 + \exp(\Delta G_{mix}(Q)/\beta_{DFT}))$. The smoothing parameter β_{DFT} is calibrated so that ΔG_{mix} values approximately 0.025-0.05 eV/f.u. above the convex-hull threshold correspond to

intermediate probabilities ($\approx 70-90\%$), thereby embedding the known DFT uncertainty directly into the constraint. The probabilistic stability prior multiplies the Expected Improvement acquisition function, penalising, but not strictly excluding, compositions that are likely unstable. In this way, the Bayesian Optimisation loop remains robust to uncertainty in the theoretical constraints and avoids brittle, overly sharp decision boundaries.

An instability index I_c , defined as the integrated colour change of an encapsulated perovskite film over an accelerated degradation test duration T , was chosen as the figure of merit for optimizing stability (objective function) and measured experimentally. Expected Improvement was used as acquisition function and the surrogate model was based on a Gaussian Process.

1.3.2 Gaussian Processes and Bayesian Optimisation in Graphene Research

There are several examples of the use of GPs related to graphene, covering both simulations and experimental work. Zheng *et al.* developed a GP metamodel to predict the density, Young Modulus and ultimate tensile strength of graphene aerogels based on the effective radius of spherical water inclusions, using data from molecular dynamics simulations [261]. Schmitz and Schnieder constructed a GP, fed with molecular simulations data, to predict atomic hydrogen adsorption on graphene using an adaptive regularisation scheme. Instead of adding a constant noise term to the covariance matrix, their approach allowed the error to be adjusted to the importance of each data point, enabling higher accuracy and numerical stability [262].

Yamawaki *et al.* demonstrated the use of Bayesian Optimisation to solve a typical dilemma of conflicting demands in thermoelectric materials, where high electrical conductivity and Seebeck coefficient but low thermal conductivity is desirable [263]. The machine-learning approach was used to design and evaluate, using first-principles calculations, an optimal structure for graphene nanoribbons with an efficiency more than five times that achieved by random search. Figure 1-8 compares the performance of Bayesian optimisation against random search when identifying graphene-nanoribbon structures with high thermoelectric efficiency. The vertical axis represents the number of evaluations required to find a candidate within the top fraction of performers, normalised by the number required under random search. Values below unity therefore indicate a computational speed-up. As illustrated in the figureFigure 1-8, Bayesian search accelerates the exploration in all cases, except for a singular point for the best case of

$m=6$ (m is the number of vacancies in the graphene nanoribbon). They used the open-source library common Bayesian optimisation (COMBO) [264].

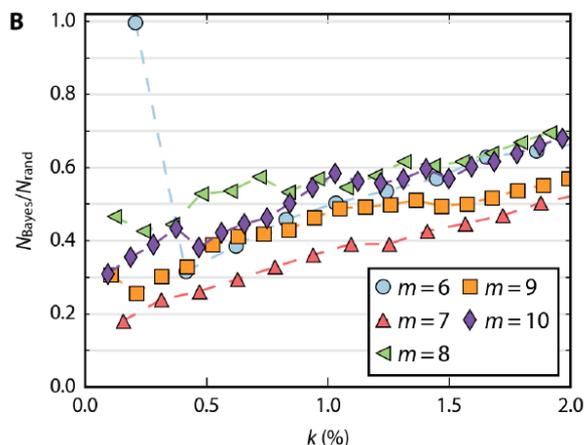


Figure 1-8. Efficiency of the Bayesian optimisation. $N_{\text{Bayes}/\text{rand}}$ is the average number of calculations needed until, for the first time, a structure that belongs to the top $k\%$ of all the candidates is found with the Bayesian/random search. Reproduced with permissions from [263].

The authors used a Bayesian linear regression model with an objective function consisting of a figure of merit that contains the parameters that affect thermoelectric performance, and the descriptors of the graphene nanoribbons as \mathbf{x} vector. Prior distribution was obtained using random feature map while the next point for the observation was determined using Thompson sampling.

Although these studies showcase sophisticated use of GPs and BO, they all rely exclusively on theoretical or simulation data. While this is valid in the context of exploring design spaces where the data are noise-free and fully controlled, one of the unique advantages of GPs is their principled treatment of uncertainty, which becomes particularly valuable when dealing with noisy experimental measurements. Moreover, BO is especially impactful when the number of experiments is severely constrained by cost or time and the choice of each new experiment must be informed by uncertainty in previous measurements, a regime that is less critical in purely computational studies where data generation is comparatively inexpensive and dense sampling is generally feasible.

Chaney *et al.* utilized GP to model the response of three target variables: productivity yield, conversion rate of graphite to graphene and conductivity of the graphene films, to the loadings of graphite and ethyl cellulose in a wet jet milling process [265]. BO was used, resulting in the stabilisation of the Pareto frontier after four batches of five tests each, from an initial dataset of nine points. Given the low-dimensional nature of the design space and the use of a structured initial experimental design, rapid convergence is not unexpected. However, the key contribution of this work lies in demonstrating that

BO can efficiently organise experimental sampling and avoid redundant testing under conditions where each experiment is costly and time-consuming, achieving performance comparable to conventional liquid-phase exfoliation studies that typically require similar numbers of experiments but explore much narrower parameter ranges without a systematic optimisation framework.

Patel *et al.* [266] used BO to find the best combination of graphene, Kevlar aramid nanofibers and carbon nanotubes to maximize the performance of supercapacitor electrodes. Such performance was computed using a utility function (objective function) which combined electrochemical features (specific capacitance at different scan rates) with mechanical properties (Young's modulus, tensile strength, toughness and ultimate tensile strength). The importance of electrochemical versus mechanical properties in the utility function is controlled by the parameter a , which serves as a weighting coefficient in the performance metric $U = (1 - a)ECU + (a)MU$. Here ECU (electrochemical utility) reflects normalised capacitance values measured at multiple scan rates, while MU (mechanical utility) combines normalised Young's modulus, ultimate tensile strength, toughness, and strain at break. Thus, a determines how strongly the overall utility favours mechanical performance over electrochemical performance. In the article, the authors set $a = 0.5$, meaning that both contributions were weighted equally. GP regression was used to fit the initial set of experimental data, providing a prior distribution. Expected Improvement was chosen as acquisition function.

The models were obtained from the free Machine Learning library Scikit-learn, which is built on Python [267]. The first iteration suggested by the model yielded a composition that had a 5.5% improvement in utility over the best performing composite (92.625 wt% rGO, 4.875 wt% ANFs, and 2.5 wt% CNTs) from the initial experimental data. The new composition outperformed the previous best performing electrode with an increase of 78.8% and 34.0% of Young's modulus and strength, respectively. However, the capacitance at $1 \text{ mV}\cdot\text{s}^{-1}$ was $117.2 \text{ F}\cdot\text{g}^{-1}$, representing a 29.6% decrease. This is because the objective function equally weights mechanical and electrochemical performance while the initial experimental data was biased towards high specific capacitance values corresponding to high rGO loadings. According to the authors, the area from the design space suggested by the model would not have been explored as rapidly when using traditional experimental approaches, based on scientist intuition. However, no direct benchmarking against alternative model-based or classical DoE strategies was performed to quantitatively support this claim.

These works are important experimental demonstrations of BO, but they focus on mixture proportions or additive compositions rather than synthesis parameters. This means they optimise the formulation of graphene-containing systems rather than controlling the intrinsic properties of the graphene itself through processing conditions. As a result, there remains a gap in understanding how synthesis variables directly impact material performance.

Alghfeli and Fisher used a GP with Matérn 5/2 kernel as surrogate model combined with Bayesian Optimisation to obtain the best reaction conditions to synthesize graphene sheets by Chemical Vapor Deposition with minimal I_D/I_G ratio [268].

Wahab *et al.* [269] used BO to support the production of laser-induced graphene. The ratio of G and D peaks from Raman spectra, which are characteristic bands used to assess graphitic order and defect density (see 2.2.7 for details), was taken as objective function, while laser power, irradiation time, gas pressure and gas type were the parameters forming the \mathbf{x} vector. Random forest with 500 trees was used to build the surrogate model: this is compatible with an \mathbf{x} vector containing both continuous numerical and discrete categorical parameters, as in this case. Expected Improvement is used as acquisition function. Thanks to this approach and after 50 optimisation iterations, the authors claimed to have obtained laser-induced graphene with Raman G/D ratios a factor of four larger than those found in the literature. The library used was mlrMBO [270].

Although the optimisation problem involved only four input parameters, Wahab *et al.* employed 50 Bayesian optimisation iterations because the effective search space was extremely large, heterogeneous and experimentally noisy. The parameters span hundreds to hundreds of thousands of possible values and include both continuous and categorical variables, motivating the use of a random-forest surrogate model. Moreover, the authors explicitly acknowledge variability arising from substrate effects and batch-to-batch differences in graphene oxide films, which necessitates iterative surrogate model refinement. Finally, the optimisation budget was dictated by experimental constraints rather than dimensionality alone: the authors explicitly fixed a total budget of 70 experiments (20 initial random samples plus 50 BO iterations), corresponding to the maximum number of reliable laser patterning experiments achievable within a work week.

While these two studies focus on synthesis parameters, its optimisation target is limited to structural metrics (Raman ratios) rather than functional electrochemical properties.

For energy storage applications, a direct link between synthesis conditions and electrochemical performance would be more informative.

Ebikade *et al.* [271] used BO to perform data-driven experimental design for synthesizing graphene-based structures as catalysts for the Hydrogen Evolution Reaction. The synthesis parameters space consisted of three variables: final temperature, heating rate and hold time. Altogether form the input parameters, or \mathbf{x} vector. The objective function to maximize is the total nitrogen content in the graphene sample corresponding to possible values of \mathbf{x} .

Latin hypercube design was used to get the initial 10 samples of the three-dimensional experimental space. Normally distributed multivariate function was used as surrogate model (GP) and Expected Improvement as acquisition function. According to the authors, a total of 40 experiments would have been needed to gain similar or even less knowledge about the experimental space than the one gained using Bayesian optimisation, which required less than 20 observations. This comparison was made within a single Bayesian optimisation implementation using the open-source Python package pyKriging [272], and no evaluation of alternative Gaussian Process models or software platforms was reported.

While Ebikade *et al.* clearly demonstrate the efficiency of BO for experimental design in electrochemical applications, their study is limited to the fact that the material under investigation is activated carbon rather than graphene. Consequently, although their results illustrate how BO can accelerate the exploration of synthesis-condition space and reduce experimental costs, the direct relevance to graphene remains uncertain given the differences in structure and properties between the two materials. This suggests that extending such methodologies to graphene-specific system is an important next step, and it is in this direction that the present work is focused. On the other hand, while existing studies have demonstrated the potential of GP and BO for guiding the synthesis of graphene-based materials, these approaches have predominantly treated the system as a purely black-box problem, relying solely on input-output data relationships. However, real-world experimental processes, particularly in complex systems (e.g., the reduction reaction of graphene oxide) often contain partial prior knowledge, such as known sources of noise or physically informed constraints, that can be integrated into the modelling process. This motivates a shift toward grey-box modelling, which lies between the extremes of black-box and white-box models and which are gaining attention in the literature [273,274]. Grey-box models combine the interpretability and structure of white-box models with the flexibility of black-box models, making them

particularly useful in domains where partial domain knowledge is available, but the system is too complex to be fully described analytically.

Among these, physics-informed GPs are examples of grey-box models which are emerging as a powerful approach to integrating physical laws or constraints into the GP framework. Unlike traditional GPs that rely solely on data, physics-informed GPs incorporate governing equations (e.g., differential equations) or physical principles as prior information. This allows the model to respect the underlying physics/chemistry/biology while still leveraging the flexibility of data-driven learning. As a result, these models are particularly effective in scenarios where data is sparse or noisy, as the physical constraints guide the model toward more plausible and interpretable solutions.

A notable example of a physics-informed GP in supercapacitors is the work of Pan *et al.* [275]. The model relates capacitance to pore characteristics (micropore surface area, S_{micro} , mesopore surface area, S_{meso}) and scan rate (ν). In the physics-informed formulation, the GP does not directly predict the capacitance; instead, it is trained to predict the transformed output $y = \ln(C_0) - k\nu$, derived from a physically motivated rate-dependent capacitance model. The capacitance is subsequently recovered as $C = \exp(y)$. The authors compared this approach with a fully data-driven artificial neural network (ANN) and a conventional, non-physics-informed GP. Although the physics-informed GP exhibited slightly higher RMSE during cross-validation, it preserved physically consistent electrochemical behaviour at high scan rates, correctly predicting a monotonic decrease in capacitance, whereas the ANN yielded unphysical negative values and the conventional GP predicted increasing capacitance at high scan rates. This demonstrates that the physics-informed model provides predictions that remain aligned with known physical laws and mechanisms, unlike the purely data-driven approaches.

1.3.3 Machine Learning Models and Explainability of Graphene Capacitance

Given sufficient data and a rule-discovery algorithm, a computer can employ Machine Learning to uncover both known and previously unknown physical laws without human intervention [276]. Nevertheless, such data-driven discovery approaches are not without limitations. ML models may infer correlations that do not reflect true underlying causal mechanisms, particularly when the dataset is incomplete, biased, or noisy. This can lead to the generation of artefactual relationships or 'laws' that reproduce the training data but fail to generalize. Moreover, symbolic-regression-based methods or other rule-extraction techniques may overfit, producing mathematically elegant expressions that lack physical

validity outside the sampled domain. Consequently, ML-derived physical rules must be critically assessed, validated against established theory, and ideally constrained through the incorporation of physical priors.

This capability is particularly valuable in materials research, as it enables the identification of patterns linking physical variables and material features to performance metrics obtained from experimental characterisation (e.g. corrosion resistance [277], cellular toxicity [278], elastic modulus [279], redox potential [280]). In addition, data derived from atomistic simulations can serve as a foundation for constructing new descriptors and describing novel materials within established theoretical frameworks. The potential sources of data for ML models are virtually limitless, including theoretical simulations, experimental results, scientific literature, or combinations thereof.

Schattauer *et al.* [281] used ML to construct the Hamiltonian of a graphene cell with a defect, based on atomic and electronic structures calculated via Density Functional Theory (DFT). Deterministic models like this are especially useful for explaining material properties from first principles. However, experimental realities often diverge from model predictions, as many influencing parameters may be overlooked. Conversely, fully data-driven models can achieve excellent fits to experimental data, but their complexity can make them difficult to interpret. As a result, these models may lack a clear correspondence with known physical phenomena.

In this section, a summary of studies using data-driven models to investigate graphene capacitance is presented. For those studies that address model explainability, the insights derived from the model structures are also described.

Zhu *et al.* collected data of carbonaceous materials from more than 300 publications, generating a dataset with 681 instances, where 5 descriptors of materials physics-chemical properties and electrochemical characterisation were chosen as input variables, and the capacitance was selected as the output [282]. After filtering which of the instances have values in all the input variables, a total of 178 of them were used to train the ML algorithms. Four different algorithms were trained, with the Artificial Neural Network achieving the best results of all, with a R^2 of 0.91. Although the accuracy of the model was clearly stated, there was no analysis about the relative importance of each of the input features in maximising the capacitance. Understanding the contribution of the features to the target variable is key to guide our experimental design of new electrodes with the objective of obtaining new materials with higher capacitance values.

Su *et al.* gathered 121 sets of carbon-based supercapacitors, with seven characteristics describing their chemical and physical properties, plus capacitance [283]. Four different

ML algorithms were used, and three error metrics were calculated to compare the results. The best results were obtained by using Regression Trees, giving the highest R^2 , and lowest Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). Moreover, by studying the structure of the trained Regression Trees, a set of rules defining the capacitance from the model input variables were extracted, concluding that the Potential Window, followed by the SSA and the I_D/I_G ratio were the factors that affected the most to the capacitance of the materials in the dataset.

Zhou *et al.* investigated the relationship between the scan rate and the surface areas of micro and mesopores in activated carbons with the power density and specific capacitance extracted from cyclic voltammograms in 6M KOH electrolyte [284]. The *Leave-one-out* cross-validation method was used to evaluate the performance of four ML models trained from a dataset composed of 70 data points. The tested models were Generalised Linear Regression, Support Vector Machine and Random Forest, Artificial Neural Networks (ANNs). K-fold cross validation with $k=10$ was used for fine-tuning, while leave-one-out cross validation was employed to extract the performance metrics of the models. To prevent overfitting in the ANNs, Bayesian Regularisation back-propagation was used, achieving the best performance in terms of both R^2 and RMSE. The ANN-trained model was used to plot heatmaps, showing the regions in the S_{micro} vs S_{meso} space with highest and lowest capacitance. From the heatmaps, it was concluded that the best results are obtained when both micropore and mesopore surface areas have moderately high values, leading to the highest capacitance. A second conclusion was that an extremely high micropore surface area can be detrimental for the capacitance, as it leads to inaccessibility for ion transport.

In a follow-up study, Zhou *et al.* trained machine learning models in which the inputs consisted of eight variables related to carbon material features, including the surface areas of micropores and mesopores, as well as the percentages of different nitrogen and oxygen doping configurations [285]. Four different machine learning algorithms were tested: generalised linear regression, support vector machine (SVM), random forest (RF), and artificial neural network (ANN). With an R^2 of 0.8495 and RMSE of 31.5491 F/g, the ANN demonstrated superior performance in predicting capacitance. According to the authors, the use of Bayesian regularisation in backpropagation was key to achieving those results. The model provided useful insights about the physiochemical conditions that favour carbon electrodes' capacitance. From the previously introduced article from the same group [284], a surface area of the micropores larger than a given threshold can lead to lower capacitance due to a reduction in electrical conductivity and increased inaccessibility of ions to the pores. Thus, a high presence of mesoporosity is

also very important, as this favours ion transport through the carbon electrode. Regarding nitrogen doping, the heatmaps extracted from the ANN model demonstrated a higher positive impact on the capacitance of the N6-doped vs both N5 and NQ-doped carbon groups, probably because of its higher contribution to pseudocapacitance coming from a lower binding energy with the potassium cation, which enhances surface activity. Oxygen doping at 1 at.% showed little effect.

Gheytaanzadeh *et al.* trained Support Vector Machine algorithms using a set of 681 carbon-based supercapacitor characterisations from 300 publications [286]. The model inputs consisted of 13 variables including physiochemical properties of the carbon structures and a parameter of the electrochemical characterisation, namely the voltage window. To find the best hyperparameters for the SVM model, a meta-heuristic optimisation technique inspired by the social hierarchy of grey wolves was applied. After an outlier analysis, only 31 points out of the 681 were categorised as such, demonstrating the efficiency of the model in widely representing the inherent relationships between the inputs and the capacitance. Sensitivity analysis showed that the Specific Surface Area (SSA) is, by far, the factor with the strongest influence on supercapacitance performance, followed by pore volume and O-doping. With an R^2 of 0.92 and RMSE of 39.2215 F/g, the authors claimed better performance compared to previously published models [282–284].

Liu *et al.* evaluated how well six different ML models explained the capacitance of porous carbon materials using 11 input features from a set of 105 samples extracted from past publications [287]. Differently from other approaches where parameters from the electrochemical characterisation were considered as input variables, in this article the authors selected a dataset where all the instances have the same Potential window, current density and electrolyte used for the characterisation of the capacitance in common. Then, only variables related to the carbon structure were considered as inputs. This choice served as their strategy to improve dataset balance by reducing variability arising from heterogeneous electrochemical testing protocols, ensuring that differences in capacitance originated primarily from material features rather than measurement inconsistencies.

It is worth noting, however, that different studies adopt different approaches to achieve dataset balance, and in this case the harmonisation of electrochemical conditions was the authors' primary method for reducing systematic variability. The best results, in terms of R^2 , RMSE and MAE, were obtained using the eXtreme Gradient Boosting (XGBoost) algorithm [288], which is based on an ensemble method that employs a regularised,

voting-mediated contribution of predictions from a set of Regression Trees. A relative importance analysis was also included, from which the ratio between the surface area of the micropores (S_{micro}) with the SSA of the material was the most important contributor to the capacitance, followed by the SSA alone.

Mishra *et al.* studied different ML models relating eight carbon physicochemical features and four parameters of the testing system with the capacitance [289]. From an initial dataset of 4,899 data entries, extracted from 147 publications, outlier extraction reduced the dataset to 4,538 instances. Then, 70% of this data was used for training five different regressors, namely Ordinary Least Square Regression, Support Vector Machine, Decision Trees, Random Forest and Extreme Gradient Boosting. The remaining 30% was used for testing, yielding the best results in terms of R^2 , RMSE and MAPE for Extreme Gradient Boosting, with 0.79, 40.27 F/g and 30.08%, respectively. This algorithm was used to train multiple sub-models, each tailored to specific values of categorical features. These features included the testing method (e.g., three-electrode vs. two-electrode), the type of electrolyte (e.g., KOH vs. H_2SO_4), and the electrode structure (e.g., activated porous carbon, hierarchical porous carbon, or heteroatom-doped carbon). Feature importance analysis was performed for each of the nine sub-models. The most frequently ranked features in the top three were SSA, nitrogen doping, pore volume, potential window, and I_D/I_G ratio. While two sub-models showed an identical top-three ranking, the remaining sub-models exhibited variations in the order of feature importance. This highlights the strong influence of categorical features on the insights derived. The authors' description of electrode structures indicates that the dataset did not include graphene or its derivatives.

Wang *et al.* trained an Artificial Neural Network (ANN) model using a dataset of activated carbon-based supercapacitor characterisations in both 1M H_2SO_4 and 6M KOH from various publications, accounting for a total of 288 data points [290]. The model inputs consisted of structural features such as micropore and mesopore surface areas, and chemical features including nitrogen and oxygen doping percentages. The dataset was split into training (70%) validation (15%) and testing (15%) subsets. The best hyperparameters for the ANN model were found using Bayesian regularisation, which made the model more robust and generalised without overfitting. This yielded a RMSE of 38.5 F/g on the test subset. This yielded an RMSE of 38.5 F/g on the test subset, which is considered acceptable given that the capacitance values in the dataset span from tens to several hundred F/g. The resulting error therefore represents only a small fraction of the overall variability and corresponds to a relative deviation typically below 10–15% for the high-capacitance samples that dominate the dataset, which is a

reasonable accuracy level of porous carbon studies based on heterogeneous data sources.

According to the ANN model, the highest specific capacitance for N/O co-doped carbon electrodes could be achieved with a micropore surface area of 1502 m²/g, mesopore surface area of 687 m²/g, nitrogen-doping of 0.5 at%, and oxygen-doping of 20 at%. Besides the model, the authors synthesized a new activated carbon using HCP-PhI polymer as precursor and applying different heat treatments under N₂ atmosphere. The carbon materials obtained demonstrated a capacitance of 610 F/g at 1 mV/s, surpassing all the values from the literature-based dataset, with a micropore surface area ($S_{micro, pores \leq 2 \text{ nm}}$) of 3650 m²/g, and a mesopore surface area ($S_{meso, pores 2-50 \text{ nm}}$) of 826 m²/g, nitrogen and oxygen doping around 1 wt% and 10 wt%, respectively. The authors re-trained the ANN using data obtained from the characterisation of the different activated carbons they synthesised, which yielded a new point of highest capacitance consisting in S_{micro} of 1710 m²/g, S_{meso} of 1050 m²/g, nitrogen-doping of 2.3 at% and oxygen-doping of 20 at%.

Although in this publication the ANN model trained with literature points achieved high accuracy levels, the addition of the new small subset from the experiments strongly affected these predictions, especially at the region of highest capacitance which was widened in the S_{micro} vs S_{meso} heatmap. This is because the authors were generating data that are outliers in both input parameters and capacitance. Thus, a robust model derived from a large set of publications may not align perfectly with new experimental data, especially when this new data exceeds the boundaries of the input and output values that define the model. This is a common limitation in machine learning models, as they perform well in interpolation tasks but are limited in their ability to extrapolate.

Across the reviewed literature, no single machine learning algorithm emerges as a universally superior approach for predicting the capacitance of carbon-based or graphene-derived supercapacitors. While artificial neural networks often report the highest predictive accuracy, particularly when combined with Bayesian regularisation, tree-based ensemble methods such as random forests and extreme gradient boosting frequently demonstrate comparable robustness, especially when trained on heterogeneous, literature-derived datasets. Overall, model performance is strongly dependent on dataset size, feature selection, preprocessing strategies, and experimental variability, which prevents the identification of a clear algorithmic 'winner'.

Moreover, most of the reviewed studies rely on purely data-driven modelling approaches. Although physically meaningful descriptors (e.g., surface area, pore size distribution,

heteroatom doping) are commonly used as input features, the learning algorithms themselves rarely incorporate explicit physical laws or electrochemical constraints. Therefore, physical consistency is not guaranteed, particularly outside the interpolation regime. Physics-guided or physics-informed machine learning approaches therefore remain largely unexplored in capacitance modelling, representing a clear opportunity for future research. Some of these models included various instances referring to graphene-based supercapacitors in their datasets.

However, most data points corresponded to amorphous forms of carbon, which may exhibit distinct properties as a supercapacitor material compared to graphene. Firstly, graphene has superior electrical conductivity, due to its sp^2 hybridised carbon atoms which facilitate charge transfer and improve rate performance [291,292]. Secondly, the tuneable surface chemistry of graphene and its derivatives, such as reduced graphene oxide, enables the boosting of pseudocapacitance through functional groups that are efficiently incorporated into the graphene network and contain redox-active atoms such as nitrogen, sulphur and/or oxygen [293]. Thirdly, graphene's excellent mechanical properties and flexibility enable the creation of more durable and flexible supercapacitor devices [294]. These combined properties make graphene-based materials promising candidates for high-performance supercapacitors, potentially overcoming the limitations of traditional amorphous carbon electrodes.

Recognising these advantages, several machine learning models have been developed to predict the capacitance of graphene-based supercapacitors. Unlike models primarily trained on amorphous carbon data, these approaches specifically incorporate graphene's unique physicochemical properties into their feature sets. Saad *et al.* [295] collected 15 features from 200 papers of graphene-based supercapacitor electrodes. The features contained physicochemical properties such as SSA, pore size or oxygen content, as well as features defining the electrochemical tests such as potential window, current density or electrolyte ionic conductivity. After filtering the outlier capacitance values, defined by the inter-quantile range (IQR), 189 remaining datapoints were used to train four ML algorithms, including K-Nearest Neighbours, Decision Trees, Bayesian Ridge and ANNs. To avoid overfitting in the ANNs, L2 regularisation was employed. Performance of the models was evaluated using R^2 and RMSE on the test split, which was not used for training the algorithms and which represented 20% of the dataset points.

The evaluation was performed using up to 5 different train/test splits, and the mean values of R^2 and RMSE were calculated as the final metrics, following the K-fold

cross-validation method with $k=5$. The best results were obtained for ANNs, with an R^2 of 0.883 and RMSE of 60.42 F/g. SHAP value analysis was performed, using the ANN model as a basis. This technique, which will be further explained in this chapter, allows for a visual representation of the relative importance of each input feature's contribution to the capacitance. The features that provided the most positive SHAP values, associated to higher capacitance, were high oxygen content firstly, followed by low carbon content, low current density, high nitrogen and high potential window.

Authors highlighted the poor performance of the models in the 4th split of the 5-fold cross-validation, which may be caused by the non-uniform characteristics of the train dataset for that split. They identified the wide range of capacitance values (15 to 910 F/g) combined with the low number of data points (189) as the potential source of non-uniformity.

A deep analysis of the dataset employed in that study reveals that many of the publications are related to graphene-metal oxide composites rather than graphene and its chemical variations. This may explain the strong influence of oxygen on the capacitance values, as materials from different works in the dataset contained metal oxide nanoparticles with proven high capacitance, such as MnCo_2O_4 [296], or nanosheets like $\text{Ni}(\text{OH})_2$ [297]. This can bias the relative importance that this feature has on graphene's inherent capacitance. In fact, graphene may be acting as a conductive support rather than an active material for providing capacitance in many of the works in the dataset. As a result, drawing conclusions about how graphene's physicochemical properties affect the capacitance from this study alone may be challenging and should be approached with caution. Additionally, the authors did not specify how missing data were handled after the filtering step. This may also be relevant in evaluating the accuracy of the model, as many of the instances in the dataset had at least one missing data point.

Chenwittayakhachon *et al.* created ML models to get insights about the effect of doping elements on graphene's capacitance [298]. Concretely, the models were trained from a dataset based on 291 data points from 17 publications. To ensure that the capacitance reaches zero for zero values of the SSA, 24 data points with zero capacitance and SSA were added to the dataset. Four different algorithms were employed: Generalised Linear Regression, Support Vector Regression, Random Forest and Artificial Neural Networks. The inputs consisted of 11 features, including electrolyte characteristics, operating conditions, structural properties and chemical composition. Like in the previously mentioned publications, ANN showed the best results, with a R^2 of 0.93 and Mean Squared Error (MSE) of 552.63 F^2/g^2 .

SHAP analysis of this model revealed that current density was the main contributor to the capacitance, which is boosted at low current densities and reduced at high currents. This is a general trend observed in GCD studies. Following this, low I_D/I_G was the second feature in order of importance that contributed to increase the capacitance, which according to the authors is likely due to the less defective and reduced doping level of the graphene. For the rest of the inputs, high oxygen levels are detrimental as well as graphitic and pyridinic-N dopings. In the case of oxygen, intermediate values of this feature (purple points) correspond to higher SHAP contribution. On the other hand, high pyrrolic-N doping and other N-dopings promote superior capacitance values.

Although the authors provide clear conclusions on the effects of varying nitrogen levels in its different chemical forms and oxygen, the limiting role of the SSA may still seem counterintuitive. Specifically, higher SSA values have a detrimental effect, which contrasts with the widely accepted notion in the literature that high SSA typically correlates with increased capacitance in supercapacitors, according to the electrical double layer theory [299–302]. Moreover, caution is needed when interpreting the high importance of I_D/I_G , as over 30% of the data for this feature was missing and imputed using the arithmetic mean. Imputing such a large amount of data can introduce bias, reduce variability, and potentially distort the model's understanding of the feature's true impact. Additionally, the reliance on mean imputation may oversimplify the data, leading to either overconfidence in the feature's role or an inaccurate estimation of its true contribution.

The same group published another work that uses a stacking strategy to build a model for the capacitance of reduced graphene oxide [303], where a total of 620 instances were extracted from 40 papers, creating a dataset of 7 inputs, with the capacitance as output.

The dataset contained data from characterisation in both KOH and H_2SO_4 electrolytes. In fact, the electrolyte concentration was one of the inputs. The architecture of the stacking model consisted of a base model which could be up to 5 different variations of tree-based methods, and a top layer of either polynomial, SVM with linear kernel or ridge regression. The dataset contained missing values for some of the features, which were imputed using KNN, approximating the missing value by the value of its nearest neighbours. The training, evaluation and selection of the stacked model architecture was carried out by combining Montecarlo Cross-Validation with K-Fold CV with $k=10$. The best model achieved an MAE of 14.3 F/g. The best model achieved an MAE of 14.3 F/g. Although MAE gains meaning only through comparison, this value is relevant because, according to the authors, it represents the lowest error across all 865 evaluated models,

outperforming the best standalone model (14.8 F/g). Moreover, they note that an error of ~14 F/g is comparable to the experimental variability typically reported for capacitance measurements of graphene-based supercapacitors, indicating that the model achieves a level of accuracy that is practically consistent with real-world electrochemical variability. While this was not explicitly stated by the authors, the SHAP analysis showed that SSA was the most important physicochemical feature, followed by %Nitrogen and %Oxygen. In a recent work from the same group [304] with a much larger dataset consisting of 1089 instances, the top 3 features were the same, but in this case, nitrogen ranked first, followed by SSA and oxygen concentration. The authors compared their results with other works, analysing the different sources of inconsistencies with other works conclusions such as differences in dataset size or in the distributions of data for the same feature across different datasets.

Kaya *et al.* trained six Machine Learning models using a dataset based on 1723 instances from 94 articles, with capacitance as objective variable and up to 12 input variables [305]. The input variables included both numerical and categorical aspects describing the material, its synthesis condition and the electrochemical characterisation. From the six types of algorithms, artificial neural network obtained the best results, with RMSE of 26 F/g. The study was based on capacitances measured using cyclic voltammograms in various aqueous electrolytes, including graphene materials together with activated carbon, carbides, nanotubes and xerogels. The total number of instances referring solely to 'graphene' was 174, of which 11 samples had oxygen levels exceeding 10%. Consequently, compared to other forms of carbon, the contribution of reduced graphene oxide electrodes to the trained model dataset is limited.

Although ANNs consistently appear as the best-performing models across the reviewed works, this superiority should be interpreted with caution. A common denominator among many of these studies is that the reported validation strategies are not strictly designed to guarantee independence between training and testing data. As a result, the performance metrics may be inflated due to various forms of data leakage or unclear separation between hyperparameter tuning and final evaluation.

For instance, Zhou *et al.* [284] combined k-fold cross-validation and leave-one-out evaluation, which compromises the independence of the test set; Su *et al.* [283] reported a 10-fold procedure without reserving an external test set, so the metrics correspond only to internal validation; Zhou *et al.* [285] first applied k-fold cross-validation (k=17) on the entire dataset and then reported a 70/15/15 split for training, validation and testing, but the arithmetic of this partitioning is inconsistent with the dataset size (N=170), raising

questions about the true separation of test data; and Saad *et al.* [295] reported both a fixed 80/20 split and additional k-fold validation, which leaves the exact evaluation protocol unclear.

Even in cases where the splitting procedure is more rigorous, meaning that the authors explicitly separated the data into training, validation, and test subsets with a clear external test set (e.g., Chenwittayakhachon *et al.* [298], Kaya *et al.* [305]), the random allocation of datapoints can still result in training and test sets containing highly similar samples from the same publication, since many instances correspond to nearly identical materials characterized under different electrochemical conditions. Under such circumstances, high-capacity models such as ANNs are particularly prone to capturing these intra-paper regularities, which artificially boosts their predictive accuracy compared to more parsimonious algorithms.

Furthermore, the use of relatively complex ANN architectures on small datasets (e.g. Saad *et al.* [295] with only 189 instances) raises additional concerns about overfitting despite the application of regularisation. Taken together, the apparent dominance of ANNs in capacitance prediction may reflect, at least in part, methodological limitations rather than an intrinsic superiority of the algorithm. Table 1-7 below summarises the main characteristics of the best-performing models that include graphene in their dataset.

Table 1-7. Summary of ML models predicting capacitance based on characteristics of carbonaceous material, including graphene.

| No. publications | Dataset size | Model Inputs | R ² | RMSE (F/g) | MAE (F/g) | CV method | Top-3 features* | Ref.** |
|------------------|--------------|--------------|----------------|------------|-----------|--------------------|---|--------|
| 41 | 121 | 7 | 0.76 | 67.62 | 52.03 | K-Fold k=10 | 1. SSA 2. Pore volume 3. Pore size | [283] |
| 300 | 681 | 13 | 0.896 | 39.22 | - | - | 1. SSA 2. Pore volume 3. % Oxygen | [286] |
| 200 | 189 | 15 | 0.88 | 60.42 | - | K-Fold k=5 | 1. % Oxygen 2. % Carbon 3. % Nitrogen | [295] |
| 17 | 315 | 11 | 0.926 | 23.51 | - | K-Fold k=10 | 1. I _b /I ₀ 2. % Oxygen 3. % Graphitic Nitrogen | [298] |
| 40 | 620 | 7 | - | - | 14.3 | Montecarlo s=10 | 1. SSA 2. % Nitrogen 3. % Oxygen | [303] |
| 91 | 1089 | 9 | 0.559 | 58.0 | 37.9 | Montecarlo s=5 | 1. % Nitrogen 2. SSA 3. % Oxygen | [304] |
| 94 | 1723 | 16 | - | 26.5 | 18.32 | K-Fold k=10 | 1. SSA 2. % Nitrogen 3. Pore volume | [305] |

*according to the feature importance analysis described in the corresponding publication. Features related to material synthesis and/or electrochemical characterisation conditions were excluded from the ranking.

**Zhu *et al.* [282] was excluded due to the lack of feature importance analysis in the publication.

While the table compares models based on similar metrics, it is important to note that the underlying datasets are not homogeneous. Each model utilises datasets composed

of different sources and inputs. This heterogeneity implies that direct, quantitative comparisons between the reported performances should be approached with caution. However, the comparison remains valuable because it allows the identification of broader trends that emerge across studies. Although all studies describe their materials as graphene-based, this apparent uniformity masks relevant differences in synthesis, composition and dataset construction. These variations introduce substantial variability that must be considered when interpreting and comparing their reported results.

For example, [283] and [298] included capacitance characterisations obtained with KOH and H₂SO₄ electrolytes in their works, whereas [305] and [295] also incorporated data using other electrolytes like Na₂SO₄. In fact, excluding the model used in [283], the rest of the models incorporate characteristics of the electrolyte (concentration, chemistry, ion-mobility...) as inputs of the model. Features related to electrochemical characterisation techniques are also used as inputs, with [295] and [298] including current densities from GCD curves, and [305], incorporating scan rates from CV curves. However, the models in refs. [286] and [283] exclude these variables. In the latter case, the authors mention a restriction in the dataset to points characterised at low current densities (≤ 1 A/g).

Despite these differences, a pattern does emerge when considering the ranking of physicochemical features. Across the majority of studies, models consistently identify specific surface areas as one of the strongest predictors of capacitance. Oxygen appears in the top three for five models. Pore volume is featured in the top three for three models, with two second-place rankings and one third-place ranking. Nitrogen ranks third twice, second twice and first once. The convergence of these findings suggests that, although the datasets and modelling approaches vary, the models tend to attribute predictive importance to similar underlying material characteristics.

A further distinction can be made across the models, specifically the number of instances related to graphene. While in [295] and [298] the entire dataset is related to graphene materials, for [283], [286] and [305] this is not the case. In fact, an analysis of the datasets from these publications reveals that the percentage of instances containing graphene data as input are 27%, 14% and 10%, respectively. This means that in the datasets from which these models have been built, the graphene data is in a limited percentage, shared with other forms of carbon such as activated carbon or carbon nanotubes. This, combined with the previous analysis of rankings, suggests that SSA and Pore Volume are likely two of the three most important features in a general model for carbon materials. However, this may not apply when the analysis is focused specifically on graphene.

In fact, by studying the feature rankings of the models including only graphene materials, they apparently agree on the importance of oxygen and nitrogen levels for capacitance. However, a comparison of their SHAP plots (Figure 1-9 below) can lead to inconsistent conclusions. As previously stated, this type of plot shows the importance of the features and also provides information about the relationship between the feature values and the model's output. Although the underlying datasets and experimental conditions differ between studies, the comparison here focuses on the qualitative trends extracted from each model rather than on their absolute values. For instance, in [295] both high oxygen and nitrogen content result in high capacitance, whereas in [298] the highest oxygen levels are detrimental, along with high levels of graphitic-nitrogen. Furthermore, while avoiding low I_D/I_G ratios is one of the most important physicochemical characteristics recommended by the [298] model to achieve high capacitance, this feature has the least relevance in the [295] model.

Taken together, these inconsistencies indicate that, beyond algorithmic performance, a critical methodological issue emerges across many of the reviewed studies: the robustness of the reported validation strategies.

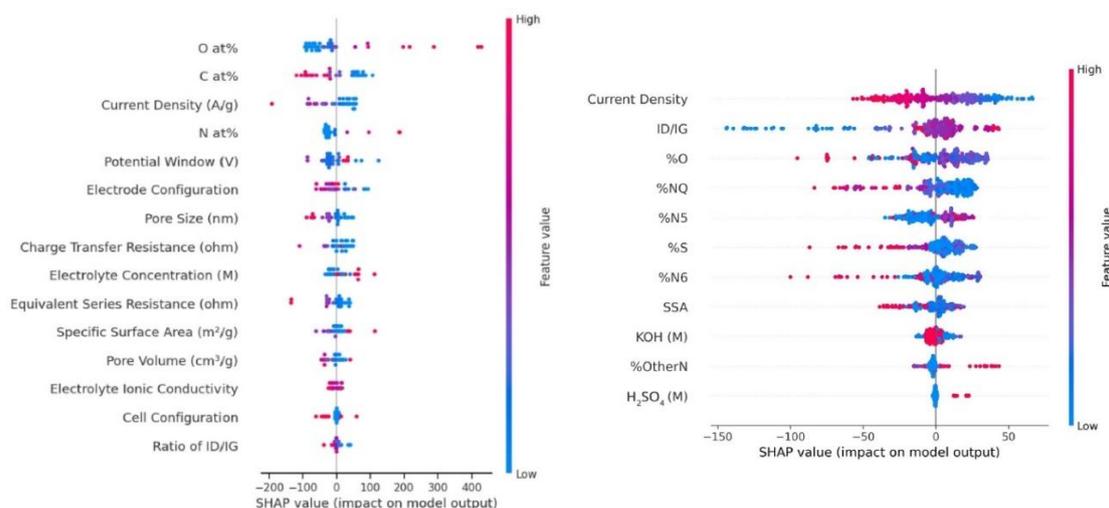


Figure 1-9. SHAP plots of ML models for graphene capacitance: (left) [295], (right) [298], reproduced with permission.

1.3.4 Research Gap

The literature presented here reveals a clear research gap concerning the optimisation of graphene oxide reduction processes through ML. While several studies have examined how synthesis conditions influence the properties of reduced graphene oxide (rGO), none have systematically investigated the combined effect of time and temperature on capacitance when using ascorbic acid, a non-toxic and environmentally friendly reducing agent, within the oxidation-reduction synthesis route. In addition,

although Gaussian Processes (GPs) have been increasingly used to model experimental systems, no study has yet incorporated experimentally derived noise into GP modelling in the context of graphene-based supercapacitors. This is particularly relevant given the variability typically observed in laboratory-scale synthesis. Finally, no previous work has introduced physics-informed constraints into the GP model, particularly bounds on the length-scale hyperparameters that reflect real synthesis conditions and measurement limitations. These constraints are critical to ensure that the resulting model remains applicable when scaling up the reduction process. This research gap will be tackled in Chapter 4, where a ML framework is developed to guide the chemical reduction of graphene oxide using ascorbic acid, with the aim of improving its performance as a supercapacitor electrode.

As highlighted in the literature, existing ML models for predicting the capacitance of graphene demonstrate considerable variation in both feature importance and underlying data structure. These inconsistencies arise from heterogeneous datasets that differ in electrolyte type, electrochemical characterisation methods, and material composition. In particular, models trained on mixed-carbon datasets (e.g., activated carbon, CNTs, and graphene) often generalise trends that may not accurately apply to graphene-based materials, or more specifically, to rGO. Even among models trained solely on this family of materials, feature attribution varies significantly. For example, oxygen content appears beneficial in one model yet detrimental in another. Such contradictions suggest overfitting to limited and unevenly distributed datasets, leading to misleading interpretations of the key physicochemical properties governing capacitance.

Chapter 5 addresses these challenges by developing a rGO-specific ML model that prioritises physical correctness and aims to avoid overfitting. A series of novel and domain-informed strategies are introduced during model development to ensure that the output is not only statistically reliable but also consistent with established knowledge about capacitance in this nanomaterial. This approach enables a more accurate and meaningful identification of the key physicochemical features influencing performance, helping to close the gap between data-driven modelling and experimental insight.

2 Methods and Materials

2.1 Contribution Statement

The Kish graphite material used in this study, including its expanded form and the reduced graphene oxide derived from it, was provided by ArcelorMittal Global R&D. The author performed the deposition of Pt catalyst on the reduced graphene oxide, as well as the cyclic voltammetry (CV) and linear sweep voltammetry (LSV) measurements on both the Kish-derived materials and the reduced graphene oxide samples. Interlaboratory studies carried out in the Electrochemical Innovation Lab at UCL were designed and performed by the author, supported by UCL student Zahra Rana.

The reduction of graphene oxide from external graphene oxide samples, along with the associated capacitance studies, was coordinated and supervised by the author and executed by laboratory technicians. Additional interlaboratory studies conducted in the Department of Physics at UCL were coordinated and partially performed by the author, designed by Dr. Srinivas Gadipelli and supported by UCL student Hanieh Akbari.

Gas physisorption analyses were performed by technicians at the ArcelorMittal Global R&D laboratories. Scanning electron microscopy (SEM), Raman spectroscopy, and elemental analysis were carried out by the external laboratory INCAR, while coordination provided by the author.

All coding machine learning work, including data extraction, curation, model development, and evaluation, was done by the author.

2.2 Experimental Techniques

2.2.1 Graphene Synthesis

Graphene can be synthesised using two general approaches: a bottom-up, in which graphene crystals are grown on a catalyst surface from a gaseous carbon such as methane; and a top-down, in which graphene layers are extracted from graphite [306]. Several strategies have been reported in the literature for isolating graphene from graphite. The most widely used approach begins with an oxidation step, which aims to increase the interlayer spacing between the graphene sheets that make up the crystalline structure of graphite. This is typically carried out using the well-known Hummers method, a wet-chemical process with many reported variations [307,308]. Following oxidation, ultrasound is applied to the graphite oxide suspension. The energy from the ultrasonic waves is absorbed by the graphite oxide particles, leading to their exfoliation into graphene oxide sheets and resulting in a stable aqueous dispersion.

The presence of oxygen-containing functional groups on graphene oxide sheets ensures the long-term stability of their dispersion in water, which is not easily achieved with water-based dispersions of pristine graphene [309]. The ability to produce stable aqueous suspensions is an important advantage when scaling up graphene production, as it greatly facilitates the industrial development of graphene-based technologies. For example, the Chinese company Sixth Element claims a production capacity of 1,000 tonnes per year [310]. Moreover, the presence of oxygen-containing functional groups on graphene oxide sheets enables further chemical functionalisation. This expands the range of potential applications, including composites, water filtration, gas separation, and other advanced material systems [311–313].

On the other hand, properties such as electrical conductivity are negatively affected by the presence of oxygen-containing functional groups. This limits the suitability of graphene oxide for applications that require efficient electron transport, such as Li-ion battery electrodes, supercapacitors, and catalyst supports in PEMFCs. Therefore, a reduction step is essential to restore conductivity and enable the use of graphene derived from graphite oxidation in electrochemical devices.

There are different means to reduce graphene oxide, using laser-based [314], biological [315], thermal [316], microwave [317,318], chemical [319], hydrothermal [320] or electrochemical [321] means. The most widely used reduction methods are chemical, employing reagents such as hydrazine [322], sodium borohydride [323], ascorbic acid [324] or plant extracts [325]. Developing reduction strategies that combine multiple approaches remains an active area of research in the graphene field. For example, in [212], a method combining graphene samples with varying oxidation levels and a two-step reduction process, consisting first of a chemical step followed by a thermal one, was shown to effectively increase the surface area of the reduced graphene oxide powder. This approach enabled the fabrication of flexible supercapacitors with record performance.

2.2.2 Graphene Oxide and Reduced Graphene Oxide from Kish Graphite

Kish Graphite is a byproduct of steelmaking which can contain up to 60% of pure graphite. Within the steel shop and its surroundings, Kish accumulates in several locations, each yielding material with varying levels of graphite purity. Figure 2-1 below shows these accumulation points at the ArcelorMittal factory in Avilés, Spain, along with the corresponding carbon content of Kish samples from each location.

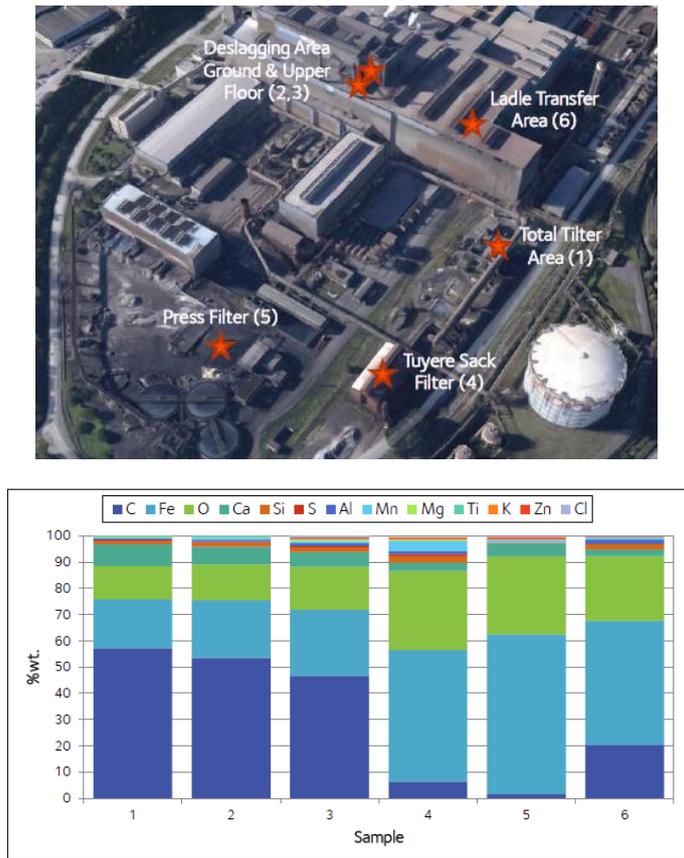


Figure 2-1. Collection spots of Kish graphite from the ArcelorMittal factory in Avilés, Spain.

Once collected, Kish graphite is sieved to separate particles based on their size. The fine fraction (<50 μm) is removed, while the coarser fraction (≥50 μm) is retained for further processing. The retained fraction is subjected to a flotation step to separate graphitic material from non-carbonaceous impurities. Following flotation, acid leaching is performed using a HCl, with an acid-to-graphite weight ratio between 0.25 and 1.0, to remove residual metallic contaminants [326].

The cleaned Kish graphite is mixed with NaNO₃ and concentrated H₂SO₄ in a mass ratio of 2:1:100 under continuous stirring in an ice bath. Then, 3.5g of KMnO₄ per gram of Kish is added gradually to oxidise the graphite, keeping the temperature below 5°C. The mixture is then transferred to a water bath at 35°C and maintained for several hours to enhance oxidation. All steps involving chemical reactions are conducted in a fume hood due to gas emissions.

Finally, the oxidised mixture is diluted with deionised water, and H₂O₂ (30%) is added to decompose residual permanganate. HCl (15%) is then used to remove manganese dioxide. The resulting graphite oxide is collected via centrifugation, washed with deionised water until neutral pH is achieved, and dried. The dried graphite oxide is redispersed in water, sonicated for exfoliation, and centrifuged to separate unexfoliated

material. The supernatant, containing well-exfoliated graphene oxide, is collected for further use.



Figure 2-2. Visual progression of material transformation. Left: Raw Kish graphite sample. Centre: Aqueous suspension of graphene oxide obtained from oxidised Kish graphite. Right: Dried graphene oxide powder after exfoliation and purification.

Graphene oxide powder is reduced following a chemical route based on the use of ascorbic acid as reducing agent. While the reduction parameters are variable, a typical procedure followed consists in adding 4 g of ascorbic acid are added into 2 L of graphene oxide dispersion with a concentration of $2 \text{ g}\cdot\text{L}^{-1}$. The reaction mixture is continuously agitated at 90°C for 3 hours to completely reduce graphene oxide sheets. A good indicator of an efficient reduction process is the accumulation of graphene at the bottom of the vessel, resulting from to the loose of oxygen-containing groups, as shown in Figure 2-3.



Figure 2-3. Suspension of ascorbic acid with reduced graphene oxide at the bottom.

2.2.3 Expansion of Kish Graphite

ArcelorMittal Global R&D holds several patents related to the synthesis of graphene-based materials from the expansion of Kish graphite [326–328]. For the results described in Section 3.2.1 of this thesis, the following expansion technique was used: a few grams of Kish graphite were added to a solution containing $(\text{NH}_4)_2\text{S}_2\text{O}_8$ and H_2SO_4 . After 5 minutes of stirring at room temperature, the expanded Kish cake was neutralised

using sodium bicarbonate and dried at 180°C for 2 hours. Finally, the powder was subjected to microwave treatment (800 W for 60 seconds), to induce further expansion.

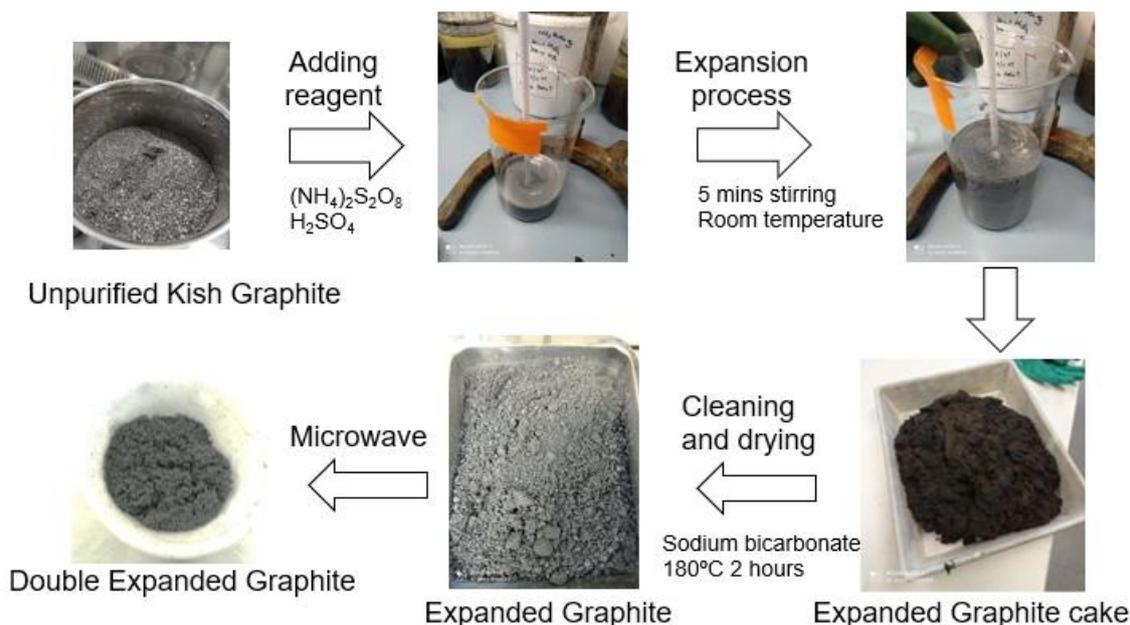


Figure 2-4. Steps for the expansion of Kish Graphite using chemical and microwave processes.

2.2.4 Cyclic Voltammetry

Cyclic voltammetry (CV) is one of the most widely used techniques for investigating electrochemical reactions occurring in a given material. It enables straightforward identification of redox potentials associated with the electroactive species under analysis. This technique provides valuable insights into the thermodynamics of redox processes, as well as the kinetics of heterogeneous electron-transfer reactions and associated electrochemical or adsorption processes.

The method involves linearly scanning the potential of the working electrode using a triangular waveform. The potential is swept from an initial value (V_1) to a final value (V_2) at a defined scan rate ($\text{mV}\cdot\text{s}^{-1}$). If only a single sweep is performed from V_1 to V_2 without reversal, the technique is referred to as linear sweep voltammetry. The resulting plot of current versus applied potential is known as a cyclic voltammogram (CV).

To perform CV measurements, the experimental setup includes an electrochemical cell, a potentiostat, a current-to-voltage converter, and a data acquisition system [329]. Figure 2-5 illustrates the basic configuration of a three-electrode electrochemical cell used for cyclic voltammetry, together with an example voltammogram obtained on a Pt electrode. The setup includes a working, reference and counter electrode immersed in an electrolyte.

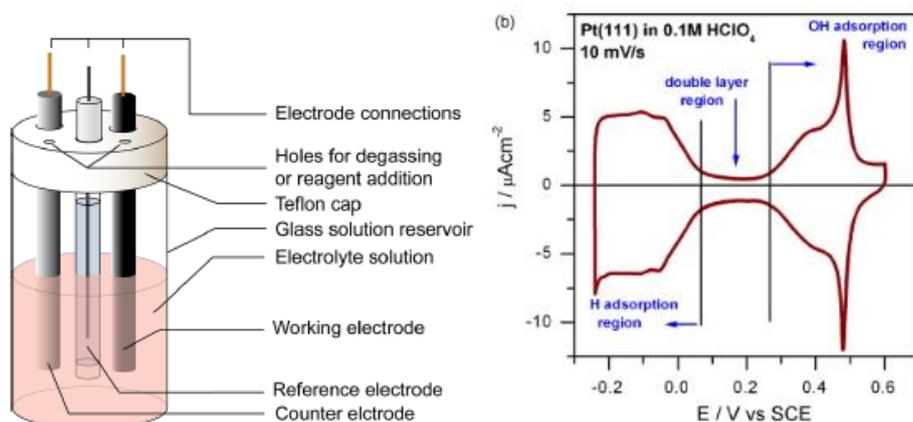


Figure 2-5. Left: Schematic representation of an electrochemical cell for CV experiments. Reproduced from [330]. Right: Cyclic voltammogram of polycrystalline Pt (1 1 1) in 0.5 M HClO₄ at 10 mV/s, from [331].

The potential of the working electrode varies linearly with time during a cyclic voltammetry experiment, while the reference electrode maintains a constant potential. Common reference electrodes include the saturated calomel electrode (Hg₂Cl₂) and the silver/silver chloride electrode (Ag/AgCl). The counter electrode completes the electrical circuit by conducting current between the power source and the working electrode. The electrolyte supplies ions to the system, enabling oxidation and reduction reactions to occur at the electrode interfaces. In the case of metallic electrodes in aqueous electrolytes without pronounced redox activity, cyclic voltammograms commonly exhibit three distinct regions:

1. The hydrogen region, at low potentials, where current peaks result from the adsorption and desorption of hydrogen on the material under investigation.
2. The oxygen region, at high potentials, where Faradaic reactions occur. In this region, the material is either oxidised (positive currents, following the IUPAC convention [330]) or reduced (negative currents), depending on the scan direction.
3. The double-layer region, situated between the hydrogen and oxygen regions, where no current peaks are observed. This indicates the absence of Faradaic reactions, and the current is primarily due to the charging and discharging of the electrochemical double layer.

The double-layer region provides insight into the electrochemical surface area of the catalyst. By contrast, cyclic voltammograms of typical supercapacitors electrode materials are typically dominated by this double-layer region, since their charge storage is largely capacitive and does not involve hydrogen or oxygen regions associated with adsorption or faradaic processes. However, materials such as reduced graphene oxide

can exhibit pseudocapacitive contributions, introducing faradic redox processes into the CV, as will be discussed later in this thesis.

By integrating the current with respect to time within this region, the total electrostatic charge associated with the formation of the Helmholtz double layer at the electrode surface can be determined. This charge, when divided by the width of the applied voltage window, allows for the estimation of the double-layer capacitance. This relationship can be described mathematically as:

$$C = \frac{i(t)}{dv(t)/dt} = \frac{Q}{V_2 - V_1} \quad \text{Equation 2-1}$$

where C is the capacitance, $i(t)$ is the measured current as a function of time, $dv(t)/dt$ is the voltage sweep rate, Q is the total charge stored and $V_2 - V_1$ is the applied voltage window.

To obtain a reliable estimation of the capacitance, the electrochemical cell should be preconditioned by cycling it for at least 20 cycles before recording the data. These cycles should be performed within a potential range that avoids the onset of redox reactions, typically from 0 V to the maximum voltage. The maximum voltage should correspond to the electrochemical window of the electrolyte, similar to those used in commercial cells. For aqueous electrolytes, this range is generally from 0 to 1 V, while for organic electrolytes it extends from 0 to approximately 2.5–2.7 V [332].

2.2.5 Linear Sweep Voltammetry and Rotating Disk Electrode

Linear sweep voltammetry (LSV) involves scanning the potential from a lower limit to an upper limit using a linear potential waveform, where the slope is determined by the scan rate. An LSV curve is generated by plotting the measured current as a function of the applied potential.

A rotating disk electrode is a device used to establish a well-defined hydrodynamic flow that enhances mass transport primarily through convection. It creates a stable, laminar flow of the electrolyte toward the electrode surface. The RDE consists of a polished, conductive disk that serves as the working electrode, embedded in an insulating cylindrical body. The entire assembly is rotated around an axis perpendicular to the electrode surface [333]. In a linear sweep voltammetry (LSV) experiment using a RDE, the angular rotation rate, ω (rad/s), determines the mass transport-limited current, as described by the Levich equation:

$$i_d = B\omega^{\frac{1}{2}} \quad \text{Equation 2-2}$$

where B is the Levich constant, defined as:

$$B \equiv 0.62 n F A D^{2/3} \nu^{-1/6} C \quad \text{Equation 2-3}$$

With n the number of electrons transferred in the half reaction, F the faraday constant, A the electrode area, D the diffusion coefficient, ν the kinematic viscosity of the electrolyte and C the analyte concentration. To model the electrical current at an electrode from an electrochemical reaction, the Koutecký-Levich equation is applied:

$$\frac{1}{i_m} = \frac{1}{i_k} + \frac{1}{i_d} \quad \text{Equation 2-4}$$

Thus, in a rotating disk electrode experiment, this expression is simplified as follows:

$$\frac{1}{i_m} = \frac{1}{i_k} + \frac{1}{B\omega^{1/2}} \quad \text{Equation 2-5}$$

Equation 2-5 allows the determination of the kinetic current from the measured current, the Levich constant for the system, and the rotation rate of the working electrode. The kinetic current is a critical parameter when evaluating the performance of a material as an electrocatalyst, as it reflects the intrinsic catalytic activity. It also serves as a basis for comparing the catalytic performance of different materials.

Since the Koutecký-Levich equation accounts for both kinetic and diffusion-limited currents, RDE experiments are best suited for extracting kinetic parameters in reactions where kinetic are not overwhelmingly fast. For highly facile reactions, such as the hydrogen oxidation reaction (HOR) in acidic media, the current is almost entirely limited by mass transport, making it difficult to accurately determine kinetic parameters using RDE. In these cases, gas diffusion electrodes (GDEs) are preferred, as they enhance mass transport and enable more accurate assessment of intrinsic catalytic activity.

2.2.6 Constant Current Charge/Discharge

Constant Current Charge Discharge involves successive charging and discharging cycles of the supercapacitor or working electrode under study, using a constant current. The current is reversed when the nominal voltage V_0 (i.e., the maximum operating voltage of the device) is reached, switching the system from charging to discharging. The output of the experiment is the potential versus time curve. Figure 2-6 shows the idealised voltage-time response of a supercapacitor during galvanostatic charge-discharge cycling. Under constant current, the voltage increases or decreases linearly with time, and the capacitance can therefore be obtained from the slope of the discharge segment according to

$$C = \frac{I \cdot (t_{2V_0} - t_{V_0})}{V_0} \quad \text{Equation 2-6}$$

where:

- I is the applied current,
- V_0 is the nominal (maximum) voltage reached at the end of the charging step,
- t_{V_0} is the time at which the discharge begins, i.e., when the voltage is V_0 .
- t_{2V_0} is the time at which the voltage returns to 0 V, corresponding to the end of the discharge step.

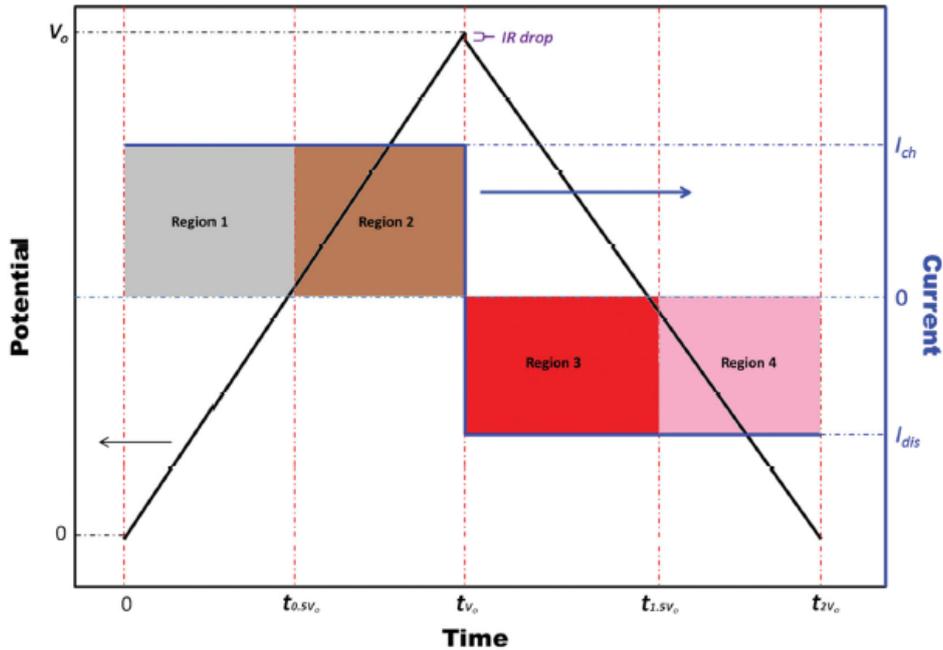


Figure 2-6. Illustration of CCCD test result from Electric Double Layer Capacitors or Pseudocapacitors with linear potential change over time. Reproduced with permission from [334].

By selecting two voltages within the linear region and measuring the time difference between them, the capacitance can be determined from the constant-current relationship $I = CdV/dt$. The schematic in Figure 2-6 also highlights different regions of the charge and discharge process, including the IR drop observed at the current reversal point, which reflects the internal resistance of the device.

CCCD test, also known as Galvanostatic Charge-Discharge (GCD), is considered one of the most versatile and accurate methods for characterizing supercapacitor (SC) devices [334]. It enables the evaluation of key performance parameters, including total capacitance, internal resistance and peak voltage. From these primary values, other important characteristics such as time constant, cycling stability, power and energy densities, leakage current and peak current, can be derived. In addition, the CCD test provides access to properties such as specific capacitance, reversibility, and the potential window of SC materials. Figure 2-7 compares the three most common configurations used for electrochemical testing of supercapacitor materials: the three-electrode setup and the symmetric and asymmetric two-electrode setups. Each

configuration distributes the applied current and measured voltage differently across the electrodes. While the three-electrode configuration isolates the response of the working electrode, the two-electrode setups reflect the full device behaviour. Understanding this distinction is essential for interpreting reported capacitance values and for comparing the results presented in this work with literature data.

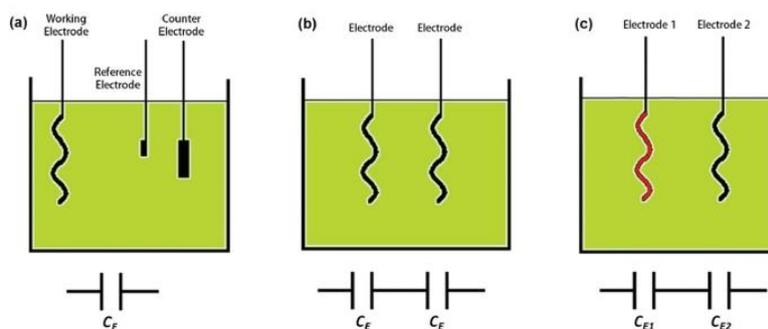


Figure 2-7. Schematic diagrams and equivalent circuits for (a) a three-electrode setup, (b) a two-electrode setup with equal electrode masses, and (c) a two-electrode setup with different electrode masses. Reproduced with permission from [334].

Three main experimental configurations are used to evaluate the capacitance of SC electrode materials: the three-electrode setup, the symmetric two-electrode setup, and the asymmetric two-electrode setup (see Figure 2-7 above). The three-electrode configuration is particularly effective for accurately determining the specific capacitance of SC materials, while the two-electrode configuration is commonly employed in SC device prototypes or final products. In this thesis, the three-electrode setup is used to evaluate the capacitance of graphene-based materials.

It is important to note that selecting an appropriate current intensity is critical for obtaining consistent and reproducible data. Other experimental conditions, such as electrode thickness, electrolyte composition and mass loadings, also strongly influence the electrochemical response. It is generally recommended [332] that the mass loading be at least $5 \text{ mg}\cdot\text{cm}^{-2}$ and the electrode thickness range between 50 and 200 μm . These constraints are important because excessively thick electrodes can introduce ion transport limitations, thereby masking the intrinsic capacitance of the active material. The applied current should be adjusted to achieve charge and discharge durations of approximately 5 to 60 seconds, which reflect the typical timescales of ultracapacitor operation.

In most applications, the electrochemical behaviour of a SC under an external load closely resembles the response of a working electrode in a CCD test [335]. Due to internal resistance, a voltage drop occurs at the beginning of the discharge curve. For

more accurate capacitance evaluation, this initial IR drop should be excluded from the voltage window used in the calculation.

Figure 2-8 below illustrates the procedure followed to prepare the electrodes for the CCCD tests. The active material, binder and conductive additive are mixed, processed into a film, cut into discs and dried before assembly into the electrochemical cell.

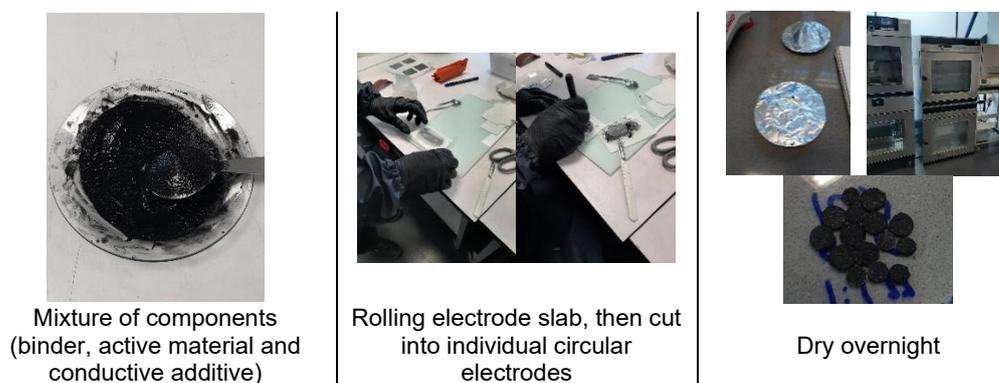


Figure 2-8. Methodology for the preparation of rGO electrodes for CCCD tests.

2.2.7 Raman Spectroscopy

Raman spectroscopy is a fast, non-destructive characterisation technique that provides insights into the structural and electronic properties of materials. It has been extensively used in the study of graphene and other carbon allotropes such as carbon nanotubes, fullerenes, diamonds, or amorphous carbon by the scientific community [336] (see left spectra in Figure 2-9 below). The key features in the Raman spectra of carbon-based materials are the so-called G (graphite) and D (defect) peaks, typically located around 1560 and 1360 cm^{-1} , respectively, for visible laser excitation. The G peak arises from the in-plane bond stretching of all pairs of sp^2 -hybridised carbon atoms in rings and chains. In contrast, the D peak originates from the breathing modes of sp^2 atoms in six-membered rings and is activated by structural defects or disorder [337].

Raman spectroscopy is highly effective in detecting structural defects in graphene, identifying functional groups, and assessing the degree of exfoliation as well as chemical and structural modifications introduced during synthesis, processing, or post-treatment. This makes it a valuable tool for quality control, enabling comparison between graphene samples from different research groups and commercial sources, and for ensuring reproducible, scalable graphene production.

In particular, the G (1560 cm^{-1}) and 2D (2700 cm^{-1}) peaks in the Raman spectra are sensitive to the number of graphene layers (see the right subplot in Figure 2-9 below). In bulk graphite, the 2D band is split into two components: 2D_1 and 2D_2 , with intensities roughly $1/4$ and $1/2$ that of the G peak, respectively. In contrast, monolayer graphene

displays a single, sharp 2D peak with an intensity approximately four times that of the G peak. This distinction makes the 2D peak a reliable indicator of graphene exfoliation and number of layers [337].

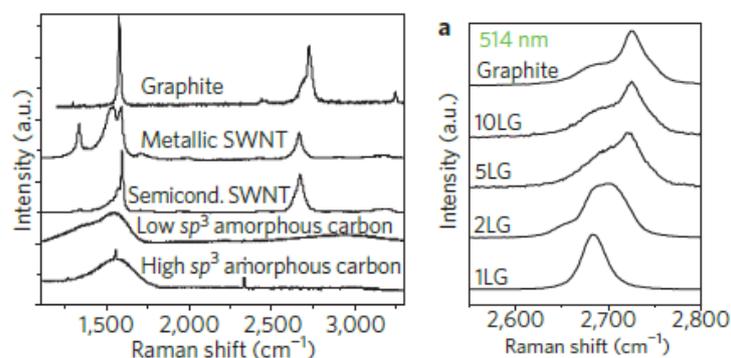


Figure 2-9. Dependence of the Raman spectra on number of layers and disorder [336]. Left: Raman spectra of graphite, metallic and semiconducting nanotubes, low and high sp^3 amorphous carbons; Right: 2D peak as a function of N for 514 nm excitation. Flakes with N layers are indicated by NLG. Reproduced with permission.

In graphene samples, Tuinstra and Koenig established that I_D/I_G varies inversely with the crystal size L_a , following the relation $L_a: I_D/I_G = C(\lambda)/L_a$, where λ is the excitation laser wavelength and $C(\lambda)$ is a wavelength-dependent proportionality constant. The coefficient $C(514 \text{ nm})$ is approximately 4.4 nm. However, this relation does not hold for amorphous carbon, where an increase in the D peak intensity corresponds to greater structural ordering, opposite to the behaviour in graphene. In graphene, edges also break translational symmetry, producing effects on the Raman spectra similar to those caused by structural defects in the sp^2 carbon network. A direct result of this symmetry breaking is the activation of the D and D' peaks [336].

Therefore, both the intensity of the D peak and the I_D/I_G ratio serve as useful indicators of the presence of edges and structural disorder in graphene sample. This is especially relevant when graphene-based materials are used as electrodes in SCs, where these factors influence the porosity, surface area, and electrical conductivity of the electrode.

In the Raman spectra of oxidised forms of graphene, the G band is typically broadened and upshifted to around 1594 cm^{-1} , while the D band at 1363 cm^{-1} becomes more intense (see Figure 2-10 below). These changes are associated with a reduction in the size of crystalline sp^2 domains caused by oxidation. In reduced graphene oxide, the Raman spectra also exhibits both G and D bands (at 1584 and 1352 cm^{-1} , respectively). However, the D/G intensity ratio is higher than that of GO. This indicates an increase in the number of sp^2 domains, but with smaller average sizes compared to those in GO [338].

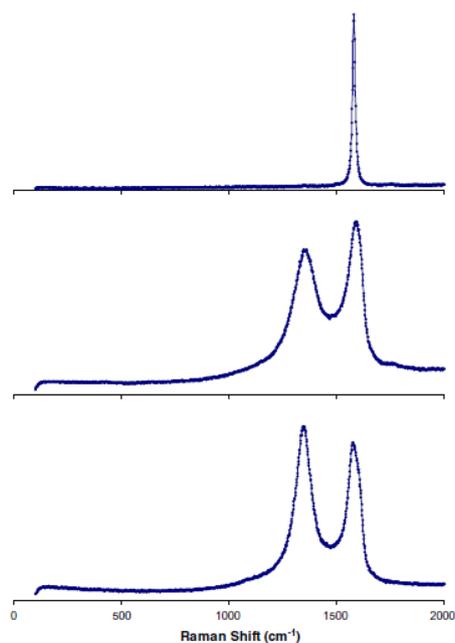


Figure 2-10. The Raman spectra of graphite (top), GO (middle), and the reduced GO (bottom).
Reproduced with permission from [338].

2.2.8 SEM

Scanning Electron Microscopy (SEM) is a widely used imaging technique that produces high-resolution images by scanning the surface of a sample with a focused beam of electrons. The interaction between the electron beam and the atoms in the sample generates signals that reveal surface features and topography. SEM is one of the most employed methods for obtaining magnified images and its extensively used across disciplines such as biology, chemistry and materials science. Furthermore, it serves as a fundamental analytical tool in the field of nanotechnology.

With respect to graphene samples, SEM is a non-invasive and effective manner for visualizing the shape and morphology of the graphene flakes [339]. While sample preparation can take time, the imaging process itself is fast, typically on the order of minutes, allowing for quick image acquisition. It is typically used to provide qualitative information about its exfoliation degree and quantitative information about the flake size distribution from a powder sample containing thousands of flakes. However, the flake population used for such analysis is restricted to the SEM's field of view. To obtain statistically meaningful results, it is recommended to measure the lateral size of at least 200 individual flakes from a given sample [340]. For more accurate quantitative analysis of graphene flake size distribution, optical microscopy, combined with specialised sample preparation, is generally preferred. Nevertheless, this thesis does not aim to quantify the flake size distribution of the graphene powders under investigation. Instead, SEM will be used to evaluate the morphological aspects of the samples.

Coating samples with gold or carbon, commonly done to reduce surface charging in low-conductivity materials, is not recommended for graphene, as it can interfere with the accurate evaluation of flake dimensions. In terms of imaging conditions, low accelerating voltages (typically below 5 kV) are preferred for graphene SEM analysis. This is because secondary electrons generated at lower voltages help minimise charging effects, and the thin nature of graphene allows the electron beam to easily penetrate the flakes, potentially resulting in a loss of surface detail [339,340].

For sample preparation for both Raman spectroscopy and SEM analysis, this thesis will follow the standardised protocols developed by the National Physical Laboratory (NPL) and the University of Manchester in 2017, which are widely regarded as the most appropriate for characterizing graphene powders [340].

2.2.9 Elemental Analysis

Elemental analysis is used to determine the carbon, oxygen, and nitrogen content in the graphene samples. The sample is heated to approximately 1200°C in an oxygen-rich atmosphere, resulting in combustion and the release of gases such as N₂, CO₂, H₂O and SO₂. These combustion products are carried by a helium flow through a gas chromatography (GC) column, where they are separated using temperature-programmed desorption (TDP).

2.2.10 Gas Physisorption

As highlighted in this study [339], only 13 out of 60 commercial samples marketed as graphene were actually graphene or few-layers graphene; the remainder were primarily graphite. This underscores the need for a rapid and reliable method to distinguish graphene powder from graphite powder, which is essential to ensure a consistent and trustworthy supply.

In this context, gas physisorption isotherms, combined with Brunauer-Emmett-Teller (BET) theory, can serve as a fast and effective method for qualitatively assessing the degree of exfoliation in graphitic powders. This approach also enables classification of various graphenic and graphitic materials based on their nanostructural characteristics. Gas physisorption offers notable advantages over microscopy-based techniques: it is faster, more repeatable, and relies on a much larger flake population for characterisation, enhancing statistical reliability.

The main challenge associated with this technique is the requirement to work with dry powder samples. During the handling of graphene dispersions, it is common for flakes to agglomerate after the drying process due to electrostatic attractive forces. This

agglomeration reduces the accessible surface area compared to a well-dispersed distribution of individual graphene flakes. Therefore, proper sample preparation is essential to ensure reproducible and comparable results.

However, as demonstrated in this study [341], this method is suitable only for qualitative, not quantitative, evaluation. This limitation arises because the specific surface area values obtained may overestimate the number of layers in the primary particles, due to the presence of agglomerates.

On the other hand, as described in section 1.3, a large specific surface area is a strong performance indicator of a material's suitability as an electrode for supercapacitors, since capacitance is given by the equation $C = \frac{\epsilon_0 \epsilon_r A}{d}$, where A represents the specific surface area.

This thesis will employ gas physisorption combined with BET theory to determine the specific surface area of graphene powders. This will serve both to evaluate the material's similarity to graphenic structures and as an early indicator of its potential performance as a supercapacitor electrode.

2.3 Machine Learning Techniques

Machine Learning is a field that combines mathematical and computational methods to automatically detect patterns in data, construct models based on these patterns, and use them to make predictions or support decision-making under uncertainty [342].

Machine Learning is usually divided into two main types:

- Unsupervised learning: where a set of N initial inputs x_i are given, forming the training dataset $\mathcal{D} = \{x_i\}_{i=1}^N$. The objective is to discover underlying structure in the data, producing outputs such as clusters, latent representations, or learned features. This process is also known as knowledge discovery.
- Supervised learning: where a labelled set of input-output pairs $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$ are given, and the goal is to establish a map between the input x_i and the output y_i variables. Depending on the nature of the output variables, supervised learning can, in turn, be divided into two categories:
 - Classification: when the outputs y_i are categorical or nominal.
 - Regression: when the outputs y_i are real-valued scalars.

Supervised learning is the most widely used form of Machine Learning (ML) in practice. This thesis focuses on regression models, aiming to establish relationships that map the synthesis conditions (input variables) of graphene-based materials to electrochemical performance indicators like capacitance or ORR onset potential.

In the context of uncertainty, probabilistic theory serves as a powerful framework for addressing questions such as: What is the best prediction of future outcomes given past data? What is the most appropriate model to explain observed data? Or, what experiment should be performed next to improve our understanding of an unknown function? Therefore, the following sections will introduce basic principles from probability theory that are widely applied in ML.

2.3.1 Parametric vs Non-Parametric Models

One of the key distinctions in supervised and unsupervised learning models is whether they employ a fixed or data-dependent number of parameters. Models with a fixed number of parameters are referred to as parametric models. Examples include linear and polynomial regression as well as neural networks, including their Bayesian variants [342]. These models are characterised by a predefined model structure, relatively efficient training, and fast prediction times.

Neural Networks (NNs) are a prominent class of parametric models. A NN algorithm performs the mapping between inputs and outputs via a set of interconnected processing

nodes distributed in layers, assigning weights and biases to the inputs [343]. Through the introduction of an activation function on each node, highly complex, non-linear data correlations can be described. Thus, NNs are good for datasets with complex problems with multiple features. The main limitation for NNs is that large amount of training data is needed to build the network, and the training process is often highly computationally intensive. Section 1.3.3 introduced several examples of NNs trained to relate the physicochemical properties of different carbon forms with their capacitance.

On the other hand, non-parametric models have no fixed parameter count, their parameters grow with the data, whereas only hyperparameters are set a priori. This means that each additional data point can increase the model's complexity. Non-parametric models are well-suited for capturing non-linear relationships and can achieve acceptable accuracy with relatively small datasets. However, they also present certain limitations. These models often perform poorly with large datasets due to higher computational requirements. In addition, their accuracy is highly sensitive to the choice of hyperparameters, which are variables that configure the model but cannot be learned directly from the data [344].

It is important to distinguish between parametric versus non-parametric models, which refers to whether the number of parameters is fixed or grows with the dataset, and deterministic versus probabilistic models, which describes whether uncertainty is explicitly modelled. For instance, Bayesian neural networks are parametric models with a fixed number of parameters while being probabilistic in nature, whereas Gaussian Process models represent a canonical example of probabilistic non-parametric approaches, as discussed later in this section.

There are two main families of non-parametric models: kernel methods and decision trees [342]. Kernel methods are a class of learning algorithms that predict the output for a new input through a linear combination of functions (kernel functions) that contain the similarity among data points within the dataset. There are multiple examples of kernels that are used in ML including Linear, Gaussian, String and Fisher kernels. Most of the kernels used are symmetric and non-negative functions. Examples of learning methods that are supported by kernel functions are Support Vector Regression, Relative Vector Machine or Gaussian Process Regression. The latter regression method will be explained in more detail below.

Decision trees can be used for both classification and regression tasks. They are constructed using simple decision rules, such as if-then-else statements, derived from the input features [345], resulting in a model that approximates the target function

through piecewise constant segments. They are easy to understand, interpret, and visualise; require minimal data preparation; and can handle both numerical and categorical features, as well as problems involving multiple outputs.

Despite their advantages, decision trees have several limitations. They are prone to overfitting, often producing overly complex models that fail to generalise well to unseen data. They are also unstable, as small changes in the training data can result in very different tree structures. Furthermore, because decision trees represent piecewise constant functions, they are neither smooth nor continuous. This limits the ability to extrapolate or make accurate predictions beyond the range of the training data [345]. To address these issues, multiple decision trees can be combined into an ensemble model to improve predictive performance. This technique is known as a Random Forest, which typically offers greater accuracy and robustness compared to individual trees [346].

2.3.2 Gaussian Processes

A Gaussian Process (GP) is a stochastic process in which any finite collection of function values follows a multivariate Gaussian distribution [347,348]. This probabilistic framework offers several advantages, including the ability to quantify uncertainty. This allows for an objective evaluation of how much is unknown about the relationship between variables and how that uncertainty is distributed across the input space. Moreover, GPs can model complex underlying patterns through the use of kernel functions and enable the use of Bayesian Optimisation (BO) in experimental design. BO is a method for efficiently finding optimal experimental conditions that has been shown to be more efficient than traditional experimental approaches, balancing exploration and exploitation through a systematic and informed approach (i.e., active learning). While BO is not restricted to GPs and can be implemented using other probabilistic models, GPs remain a particularly popular choice due to their strong uncertainty estimates and data efficiency.

A GP is fully defined by a mean function $\mu(\mathbf{x}) = \mathbb{E}[\phi|\mathbf{x}]$, which represents the expected value of the unknown function ϕ at any input \mathbf{x} , and a covariance function (or kernel), $k(\mathbf{x}, \mathbf{x}') = \text{cov}[\phi, \phi' | \mathbf{x}, \mathbf{x}'] = \mathbb{E}[(\phi - \mu(\mathbf{x}))(\phi' - \mu(\mathbf{x}'))^\top]$, which defines how function values at different inputs \mathbf{x} and \mathbf{x}' are correlated.

Training a GP involves selecting a suitable covariance function and optimising its hyperparameters, which are the adjustable parameters that govern the behaviour of the kernel. In practice, these hyperparameters are learned from data using methods such as Maximum Likelihood Estimation (MLE) or Bayesian inference. Once the kernel and its

hyperparameters are determined, the GP can be fit to the dataset. One of the most widely used kernels is the Radial Basis Function (RBF):

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^\top \mathbf{M}(\mathbf{x} - \mathbf{x}')\right) \quad \text{Equation 2-7}$$

where σ_f^2 is the variance of the function being modelled, \mathbf{M} is a symmetric, positive semi-definite matrix that encodes the characteristic length scales of the input dimensions, typically computed as:

$$\mathbf{M} = \text{diag}(l_1^{-2}, \dots, l_d^{-2}), \quad \text{Equation 2-8}$$

where $\boldsymbol{\ell} = (l_1^{-2}, \dots, l_d^{-2})$ is a vector of positive length-scales, one for each dimension. These length-scales determine the rate of decay of the exponential function, indicating how far one must move along each input dimension for the function values to become effectively uncorrelated [348]. Thus, a large length-scale implies that the covariance becomes nearly independent of corresponding input dimension, effectively making it irrelevant in the inference process. In contrast, a small length-scale indicates that the input has a stronger influence on the model, suggesting higher relevance. The RBF kernel combines high flexibility with computational simplicity, governed primarily by two hyperparameters: σ_f^2 and $\boldsymbol{\ell}$. Its high flexibility enables it to capture complex underlying patterns in the data while maintaining good generalisation performance, even with relatively small datasets.

The argument inside the exponential function essentially represents the scaled Euclidean distance between the input points \mathbf{x} and \mathbf{x}' , where the distance is divided by the length-scale. When input points are close, the covariance approaches 1, indicating high correlation. As the distance between inputs increases, the covariance decreases.

For large values of l_i , function values remain correlated over longer distances, leading to smoother functions. Conversely, small l_i values reduce the correlation range, allowing the model to capture fine-grained, rapidly changing patterns in the data. While this makes the GP capable of fitting complex functions, it also increases the risk of overfitting, particularly when the data contains experimental noise, as is common in laboratory measurements.

To mitigate overfitting, a regularisation term can be added to the covariance function in the form of an additional kernel known as the White Kernel. This kernel adds independent noise variance to the diagonal of the covariance matrix, resulting in a combined kernel of the form:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^\top \mathbf{M}(\mathbf{x} - \mathbf{x}')\right) + \sigma_n^2 \delta(\mathbf{x}, \mathbf{x}') \quad \text{Equation 2-9}$$

where the constant σ_n^2 controls the amount of noise added. This has the effect of improving numerical stability, preventing the model from fitting noise in the data, and reducing overfitting by enforcing smoother predictive functions. Once ℓ and σ_n^2 are selected, the GP provides a flexible framework for approximating a black-box function, even when starting with a small initial dataset.

In this thesis, Gaussian Processes were implemented using the Scikit-learn package [267], an open-source Python library that supports both supervised and unsupervised ML methods. Scikit-learn also provides tools for model fitting, evaluation, and data preprocessing. The Gaussian Process Regressor class [349] was used to perform regression with Gaussian Processes. By default, the prior mean is assumed to be constant and zero, although it can be customised based on the training dataset. The covariance is defined through a kernel function, whose hyperparameters are optimised by maximising the log-marginal-likelihood using a specified optimiser. The noise level in the data can be accounted for through the alpha parameter, which may either be estimated during optimisation or directly provided by the user when experimental measurements of the noise variance are available.

An illustrative example in two dimensions is shown in Figure 2-11 below. In this example, a Gaussian Process is applied using RBF kernel, and the posterior distribution is constructed after nine observations of the true underlying function. The top panel displays samples from the prior distribution, demonstrating the range of smooth functions considered plausible before any data are observed. The prior mean remains at zero, with a uniform uncertainty band across the domain. In contrast, the lower panel shows the corresponding posterior distribution. The observed data points (in red) constrain the GP, leading to a posterior mean that interpolates the observations. The uncertainty band contracts around regions with data and widens elsewhere, reflecting the model's updated confidence. Posterior samples illustrate how the GP balances adherence to the observations with uncertainty in unobserved regions.

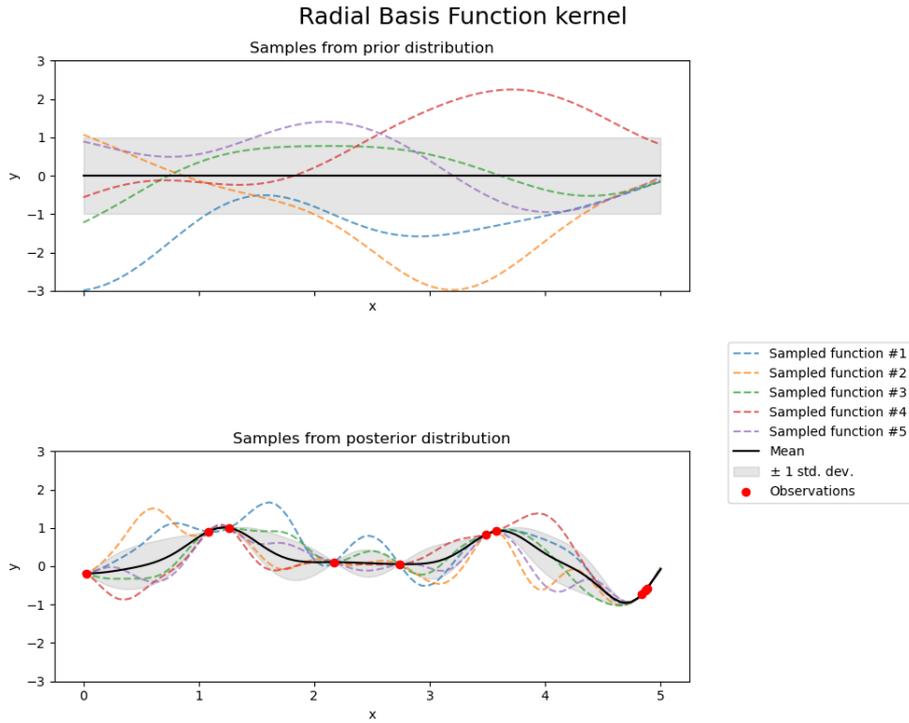


Figure 2-11. Prior and posterior distribution using Gaussian processes with radial basis function kernel after 9 observations [350].

2.3.3 Bayesian Optimisation

As previously described in this thesis, Bayesian Optimisation is a probabilistic approach for the global optimisation of a 'black-box' objective function. It works by sequentially selecting input values and evaluating the corresponding outputs of the objective function, with the aim of efficiently exploring the input space.

Mathematically, this process begins by defining a prior distribution, $p(\phi)$, over a surrogate function $\phi(\mathbf{x})$ that approximates the true objective function $f(\mathbf{x})$, where \mathbf{x} is a vector of input variables, each treated as a random variable [351]. The prior is selected based on an initial set of data points (input-output pairs), reflecting our assumptions about the likely form of the underlying function. Typically, a GP is used as the prior due to its flexibility and capacity to model uncertainty.

When a new observation $(\mathbf{x}_n, f(\mathbf{x}_n) \equiv \mathbf{y}_n)$ is obtained, Bayesian inference is applied to update the prior and obtain a posterior distribution, $p(\phi | \mathbf{x}_n, \mathbf{y}_n)$, which reflects the updated belief about the function given the new data. This update is performed using Bayes' theorem:

$$p(\phi | \mathbf{x}_n, \mathbf{y}_n) = \frac{p(\phi | \mathbf{x}_n) \cdot p(\mathbf{y}_n | \mathbf{x}_n, \phi)}{p(\mathbf{y}_n | \mathbf{x}_n)} \quad \text{Equation 2-10}$$

Where $p(\mathbf{y}_n|\mathbf{x}_n, \phi)$ is the likelihood of the new observed value, and $p(\mathbf{y}_n, \mathbf{x}_n)$ is a normalisation constant, given by:

$$p(\mathbf{y}_n|\mathbf{x}_n) = \int p(\mathbf{y}_n|\mathbf{x}_n, \phi) \cdot p(\phi|\mathbf{x}_n) d\phi \quad \text{Equation 2-11}$$

This sequential updating of the prior (i.e., the surrogate model) enables efficient approximation of the objective function using only a small number of parameters. In the case of a Gaussian model, only a mean (μ) and a variance (σ^2) are needed. Therefore, BO provides a framework for continuous refinement of our beliefs about the function as new data from experiments or simulations become available.

To evaluate the fidelity of this inductive process and guide where new observations should be taken in the experimental space, a formal Bayesian optimisation policy is required. This is achieved through the use of an acquisition function, a function that guides the selection of the next point to evaluate by quantifying how valuable each candidate input is for the optimisation process. The acquisition function assigns a score to each potential input location based on its expected contribution to the optimisation process, balancing two competing objectives:

- 1) Exploration: selecting the point that maximises the knowledge gain on the parameters space
- 2) Exploitation: choosing the observation where the objective function is apparently maximised.

There are several types of acquisition functions, each implementing a different strategy for this trade-off. Among the most widely used are the following:

- **Upper Confidence Bound (UCB):** which consists of a straightforward exploitation versus exploration trade-off, balanced by a parameter:

$$\alpha_{UCB}(\mathbf{x}; \lambda) = \mu(\mathbf{x}) + \lambda\sigma(\mathbf{x}) \quad \text{Equation 2-12}$$

Where $\mu(\mathbf{x})$ is the mean of the surrogate model as a function of \mathbf{x} and $\sigma(\mathbf{x})$ its variance; λ is a parameter that serves to weight the importance of gaining knowledge about the variance versus the mean for the next observation.

- **Probability of Improvement (PI):** computes the probability of an observed value to improve upon some chosen threshold, regardless of the magnitude of this improvement. Defining a function 'improvement' $I(\mathbf{x})$ as:

$$I(\mathbf{x}) = \max(f(\mathbf{x}) - f(\mathbf{x}^*), 0) \quad \text{Equation 2-13}$$

where $\Phi(x)$ is the Gaussian Cumulative Distribution Function and x^* is the best solution obtained so far to maximise $f(x)$. Then, Probability of Improvement can be defined as acquisition function as it follows:

$$\alpha_{PI}(\mathbf{x}; \mathbf{x}^*) = \Pr(I(\mathbf{x}) > 0) = \Phi\left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^*)}{\sigma(\mathbf{x})}\right) \quad \text{Equation 2-14}$$

- **Expected Improvement (EI):** balances the exploitation of the most stable regions and the exploration of high-uncertainty regions within the compositional space.

$$EI(\mathbf{x}) \equiv \mathbb{E}[I(\mathbf{x})] = (\mu(\mathbf{x}) - f(\mathbf{x}^*))\Phi\left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^*)}{\sigma(\mathbf{x})}\right) + \sigma(\mathbf{x})\phi\left(\frac{\mu(\mathbf{x}) - f(\mathbf{x}^*)}{\sigma(\mathbf{x})}\right)$$

Equation 2-15

where $\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right)$, i.e. the probability density function of the normal distribution $\mathcal{N}(0,1)$.

There are multiple open-source toolkits that permit the implementation of Bayesian optimisation in different forms. Table 2-1 below summarises some the most popular ones. This thesis will focus on NEXTorch.

Table 2-1. Open Source Software Libraries for Bayesian Optimisation.

| | License | Language | Year | Access, reference paper |
|-----------|--|-----------------|-------------|---|
| SMAC | Academic and Non-commercial research use | Python | 2016 | https://github.com/automl/SMAC3 |
| Spearmint | Academic and Non-commercial research use | Python | 2012 | https://github.com/HIPS/Spearmint |
| Bayesopt | GNU Affero General Public License | C++ | 2014 | https://github.com/rmcantin/bayesopt |
| Orion | BSD | Python | 2017 | https://orion.readthedocs.io/en/stable/index.html |
| BOHB | BSD-3 | Python | 2018 | https://automl.github.io/HpBandSter/build/html/index.html |
| Dragonfly | CC BY 4.0 | Python | 2020 | https://dragonfly-opt.readthedocs.io/en/master/ |
| BoTorch | MIT | Python | 2020 | https://botorch.org/ |
| NEXToch | MIT | Python | 2021 | https://nextorch.readthedocs.io/en/latest/index.html |

To summarise, experimental design can be effectively supported by machine learning to reduce the number of experiments required to achieve a specific research objective. In the context of BO, this reduction is achieved by iteratively selecting the most informative next experiments based on a surrogate model, rather than relying on exhaustive or grid-based exploration. The efficiency of this process is affected by several aspects, including the quality of data and the configuration of the surrogate model. It is also influenced by the choice of stopping criteria, such as setting a maximum number of iterations, applying expected improvement thresholds, or the surrogate-model stabilisation (i.e. assessing whether the model stops changing significantly as new data are collected). Since BO may not converge to a global optimum in complex or noisy design spaces, these choices help guide the optimisation.

The process relies on three main components: (1) a regression model trained on an initial set of known data, also referred to as a surrogate model; (2) a policy strategy that balances the trade-off between exploitation (maximising or minimising the objective) and exploration (gaining knowledge); and (3) an update mechanism based on Bayes' theorem, which incorporates newly acquired data into the model.

This cycle is executed iteratively until the research objective is met. Depending on the key performance indicators (KPIs) considered, the optimisation can be formulated either as a single-input/single-output (SISO) problem, where a single response is optimised, or

as a multi-input/multi-output (MIMO) problem, which may require a multi-output Gaussian Process.

In this work, the optimisation problem is defined by maximising the capacitance of rGO using two experimental inputs, namely the temperature and the duration of the reduction process. This corresponds to a multiple-input/single-output (MISO) formulation, which can be modelled using a single-output Gaussian Process surrogate.

The overall workflow is commonly referred to as active learning, and the complete process is illustrated in Figure 2-12 below [255].

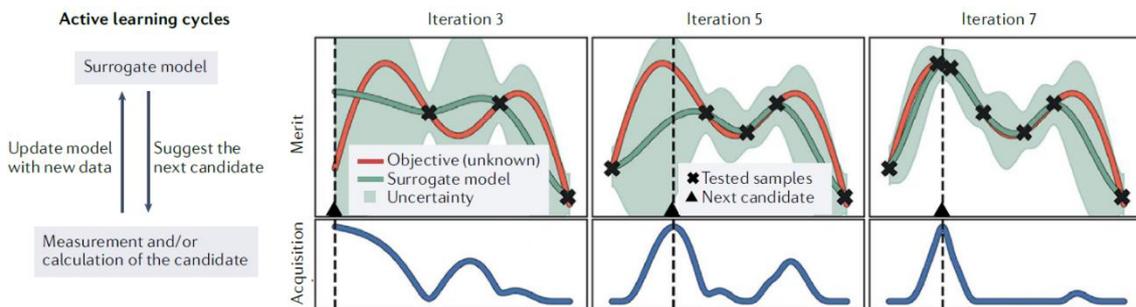


Figure 2-12. Description of an active learning cycle. Reproduced with permission from [255].

Figure 2-12 illustrates the active learning cycle used in BO. The left panel summarises the sequential workflow, which involves three main steps: (1) suggesting the next experimental candidate based on the acquisition function and the current state of the surrogate model, (2) performing the experiment or calculation to obtain the corresponding output value, and (3) updating the surrogate model with this new information. This cycle is repeated iteratively as new data become available.

The right panels show the evolution of the surrogate model over successive iterations (iterations 3, 5 and 7). In each panel, the red curve represents the true but unknown objective function, while the green curve corresponds to the surrogate model prediction, and the shaded green area indicates the model uncertainty. The black crosses denote the input points that have already been tested, and the black triangle marks the next candidate selected by the acquisition function. The lower sub-panels display the acquisition function, whose maximum identifies the next point to evaluate.

At each iteration, the surrogate model becomes progressively more accurate in regions where data have been collected, as indicated by the shrinking uncertainty band. The acquisition function balances exploration of regions with high uncertainty and exploitation of regions predicted to yield high objective values. Through this iterative process, BO directs sampling toward the most informative or promising areas of the input space, thereby reducing the number of experiments required to approach the optimum.

The active learning cycle described above illustrates how Gaussian Processes, BO, and acquisition functions operate together within an iterative framework. As the surrogate model is updated with new measurements, its predictions become progressively more accurate in sampled regions, while the acquisition function guides the exploration of promising or uncertain areas of the design space. This interplay demonstrates how machine learning can support experimental planning more efficiently than conventional trial-and-error approaches, especially when experiments are costly or datasets are sparse. Altogether, the concepts presented in this section form the methodological foundation used in the subsequent chapters, where these principles are applied to optimise the synthesis conditions and electrochemical performance of graphene-based materials.

3 Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide

Summary

This chapter explores two experimental directions motivated by the preceding literature review, focusing on performance metrics and experimental uncertainty in: (1) oxygen reduction reaction (ORR) studies of raw Kish and Kish-derived graphene, and (2) capacitance evaluation of reduced graphene oxide (rGO).

The ORR investigations are driven by iron's established catalytic activity and the unique composition of Kish graphite, a graphite-iron impurity composite derived from steelmaking. The capacitance studies are driven by emerging research on carbon-based Machine Learning models for supercapacitors and the increasing use of Bayesian optimisation for graphene synthesis.

Variability and noise characteristics are analysed across the different experimental test families. ORR reproducibility is evaluated through variability in liner sweep voltammetry (LSV), while capacitance performance is assessed via consistency in galvanostatic charge-discharge (GCD) curves.

The study introduces intentional random variability into the experimental workflow. For RDE experiments, samples were fabricated with deliberately varied drying conditions during ink deposition, simulating uncontrolled fluctuations typical of electrode preparation. Systematic variability was assessed by comparing characterisations of similar samples across different laboratories. For capacitance evaluations, rGO samples were chemically reduced, with purposeful noise introduced through variations in stirring during the reduction process.

This chapter has two primary objectives. First, to assess whether the steelmaking byproduct Kish or its derivatives exhibit potential as electroactive materials. Second, to develop a framework for uncertainty quantification in experimental tests involving graphene's synthesis and electrochemical evaluation, which will be expanded upon in the subsequent chapter discussing physics-informed Gaussian Processes for Bayesian optimisation.

3.1 Introduction

Electrochemical characterisation plays a critical role in evaluating the performance of materials for energy conversion and storage applications. While electrochemical data

are often presented with apparent precision, a deeper understanding of the sources and structure of uncertainty is essential for the reliability and reproducibility of these results.

Within this context, it is worth recalling the distinction made by Colemand and Steel in its comprehensive book [352] between measurement uncertainty, which refers to the precision (i.e. reliability and accuracy) with which a quantity is measured, and experimental uncertainty, where variability in the data arises from factors unrelated to inherent errors of the measurement system. Electrochemical measurements often involve multiple overlapping sources of experimental uncertainty, arising from factors such as electrode surface preparation, electrolyte composition, temperature stability, reference electrode drift, and instrumental limitations (e.g. potentiostat resolution, rotation speed control) [353–356].

Bhandari *et al.* [357] studied the effect of various experimental parameters on the performance evaluation of Ni-Co oxide as a catalyst for the oxygen evolution reaction (OER) under alkaline conditions. Measurements were carried out using the RDE technique, and the influence of different parameters was analysed, including the presence of Fe impurities in the electrolyte, the binder used to prepare the catalyst dispersion (PTFE vs. Nafion), and the catalyst loading. The OER activity curves were affected by all these parameters. Particularly notable was the influence of trace Fe impurities in the electrolyte, which shifted the polarisation curve to lower potentials by nearly 100 mV. In fact, the performance in the purified electrolyte, with Fe impurities below 0.08 ppb, was time-dependent, showing a decline after 2 hours compared to its performance immediately after conditioning. High mass loadings exhibited a greater tendency for bubble attachment, likely due to the increased porosity of the catalyst film. Moreover, films with low catalyst loading demonstrated poorer stability, failing to maintain performance beyond two hours, probably due to a lower availability of active sites.

The same group coordinated an interlaboratory study aimed at evaluating the influence of local laboratory conditions on ORR performance using the RDE technique [358]. A standardised measurement and conditioning protocol, including the electrochemical setup (potentiostat and electrochemical cell) and electrode materials, was implemented across all laboratories to minimise variability arising from these factors.

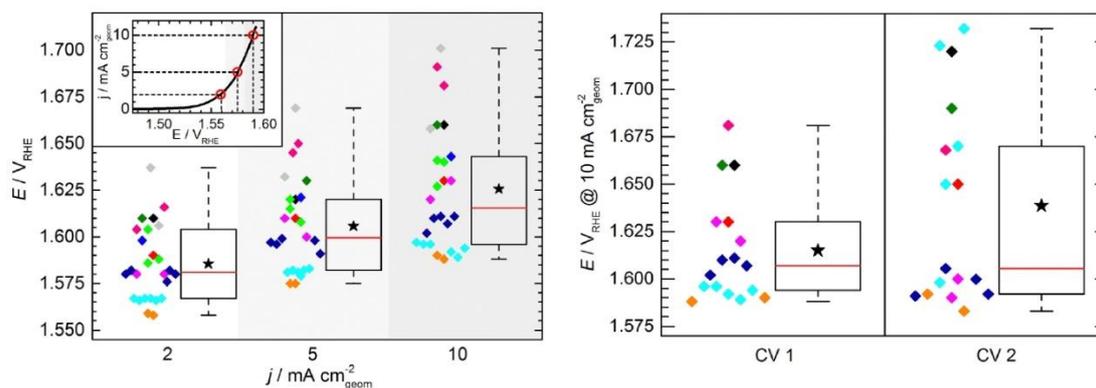


Figure 3-1. Left: distribution of potential for each of the laboratories (indicated by different colours) needed to reach 2, 5 and 10 mA cm^{-2} . Right: potential at 10 mA cm^{-2} before (CV 1) and after (CV 2) the stability measurement. Reproduced from Ref. [358] with permission from the Royal Society of Chemistry.

Figure 3-1 and Figure 3-2 illustrate the variability observed in rotating-disk electrode measurements across different laboratories. Figure 3-1 shows the distribution of potentials required to reach 2, 5, and 10 mA cm^{-2} , which were defined as key performance indicators (KPIs) by the coordinating group to enable comparison of results among laboratories. While the variability among repeated measurements within a single laboratory was less than 15 mV at 10 mA cm^{-2} , the spread increased to 110 mV when comparing results across different laboratories (see Figure 3-1), indicating that interlaboratory variability was nearly eight times higher.

The authors identified the drop-casting process used for electrode fabrication as a major source of this variability, due to its inherently limited reproducibility when performed manually. Small variations in pipetted ink volumes (on the order of a few μL) were also noted as potential sources of inaccuracy. In addition, the criteria used to judge whether a deposited catalyst film is sufficiently homogeneous for RDE testing may differ among researchers. Figure 3-2 presents photographs of catalyst films prepared by the same researcher using the drop-casting method. The images highlight the strong influence of manual deposition on the visual uniformity of the films and, consequently, on the reproducibility of the electrochemical response.

Taken together, these figures demonstrate that even when the same catalyst and nominally similar procedures are used, variations in film preparation and experimental handling introduce significant dispersion into the measured electrochemical data. This supports the conclusion that a substantial fraction of the variability reported in the literature is methodological rather than intrinsic, which strongly motivates the explicit treatment of the uncertainty associated with manual procedures adopted in this thesis.

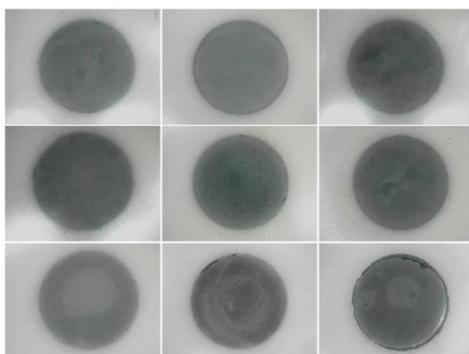


Figure 3-2. Photographs of different catalyst films deposited on glassy carbon electrodes, prepared by the same researcher. Ref. [358] with permission from the Royal Society of Chemistry.

Dos Santos *et al.* [359] analysed various factors affecting measurements of specific capacitance in materials. The authors distinguished between random errors, which can be mitigated through replication experiments and proper reporting of data along with adequate statistical analysis, and systematic errors, which are much more difficult to detect and quantify. The study identified several sources of potential errors such as temperature fluctuations, caused by inadequate thermal control of the room or poor heat dissipation by the device under test or electromagnetic fields from sources such as electric pumps, fluorescent lamps, or mobile phones near the measurement setup.

The authors further discussed how specific electrochemical techniques may introduce systematic inaccuracies if misapplied. For the CV technique, variations in scan rate influence ion adsorption in micropores, altering the double-layer formation and thus the measured capacitance, especially when the electrical series resistance (R_{ESR}) affects the CV profile. In GCD experiments, deviations from the ideal triangular profile, particularly when Faradic redox-processes are present, invalidate basic analytical equations, requiring alternative approaches to accurately assess stored and dissipated energy.

The article complemented its detailed analysis of error sources with an experimental assessment demonstrating how variations in measurement procedures and equipment can impact the reproducibility of results. Figure 3-3 below presents two tables reproduced with permission from the original publication. Table 1 (top) displays capacitance values evaluated for the same material, calculated using CV and GCD methods, under both best and worst practice conditions (a detailed description of these practices was not provided). The percent error in capacitance calculations reached 69.1% for CV and 15.7% for GCD. Table 2 (bottom) illustrates the effect of changing the potentiostat while consistently applying best practices. In this case, the measurements showed much less variability, with errors of 3.2% and 4% observed for capacitance estimations from CV and GCD, respectively.

Coleman and Steele [352] defined three categories of experiments: timewise, sample-to-sample and transient measurements. Timewise experiments involve repeated measurements under nominally steady conditions. In the context of electrochemical testing, such experiments typically require a prior conditioning step to ensure a steady state for both the material under investigation and its environment. For RDE experiments, this conditioning involves collecting CVs for several minutes under N₂ saturation to remove impurities from the catalyst surface (including oxides). In GCD experiments, a waiting period is introduced after immersing the electrodes in the electrolyte to ensure proper wetting of the pores and consistent immersion throughout the hierarchical porous structure.

Table 1
Comparisons between the best and worst practices for electrochemical characterization of SCs.

| Practice for characterization of SCs | CV ^a | GCD ^b | | EIS | | SPSC ^g |
|--------------------------------------|--|---|--------------------------------------|-------------------------------------|---------------------------------------|--------------------------|
| | C _{EDL} (F g ⁻¹) | C _{EDL} ^c (F g ⁻¹) | R _{ESR} ^d (Ω) | Z _{re} ^e (Ω) | -Z _{img} ^f (Ω) | I _{MIN} (μA) |
| Best | 9.3 | 13.1 | 12.0 | 46.5 | 1208.6 | 6.8 |
| Worst | 5.5 | 15.7 | 0.9 | 752.7 | 1966.7 | -240.0 |
| Percent error ^h | 69.1 | 16.6 | 1233.3 | 93.8 | 38.5 | 102.8 |

^aThe capacitance is obtained based on a graphical integration of the area under the voltammogram to an WVV of 1.2 V and a scan rate of 500 mV s⁻¹; ^bC_{EDL} and the resistance R_{ESR} are calculated based on the GCD discharge curve to an WVV of 1.2 V; The applied current to charge and discharge the SC is ^c(0.05 mA) and ^d(5 mA); The imaginary ^e(-Z_{img}) and real ^f(Z_{re}) part of the impedance in the Nyquist plot at a frequency of 10 mHz, and a steady state voltage of 0.3 V superimposed with a sinusoidal perturbation voltage of 10 mV peak; ^gI_{MIN} is the minimum current achieved after 60 s to a single potential step of 1.2 V (i.e., the leak current). ^hPercentage error between best and worst practice.

Table 2
Reproducibility obtained using the best practices adopted in this work for the electrochemical characterization of SCs.

| Potentiostat | CV ^a | GCD ^b | | EIS | | SPSC ^g |
|----------------------------|--|---|--------------------------------------|-------------------------------------|---------------------------------------|--------------------------|
| | C _{EDL} (F g ⁻¹) | C _{EDL} ^c (F g ⁻¹) | R _{ESR} ^d (Ω) | Z _{re} ^e (Ω) | -Z _{img} ^f (Ω) | I _{MIN} (mA) |
| BSC-810 ^h | 12.1 | 12.3 | 1.3 | 125.8 | 1040.0 | 0.22 |
| SP-200 | 12.4 | 12.6 | 1.5 | 111.4 | 979.8 | 0.24 |
| SP-150 | 12.0 | 12.1 | 2.0 | 107.3 | 997.9 | 0.28 |
| Average | 12.2 | 12.3 | 1.6 | 114.8 | 1005.9 | 0.25 |
| Percent error ⁱ | 3.2 | 4.0 | 35.0 | 14.7 | 5.8 | 21.4 |

^aThe capacitance is obtained based on a graphical integration of the area under the voltammogram to an WVV of 1.2 V and a scan rate of 500 mV s⁻¹; ^bC_{EDL} and the resistances R_{ESR} are calculated based on the GCD discharge curve to an WVV of 1.2 V; The applied current to charge and discharge the SC is ^c(0.05 mA) and ^d(5 mA); The imaginary ^e(-Z_{img}) and real ^f(Z_{re}) part of the impedance in the Nyquist plot at a frequency of 10 mHz, and a steady state voltage of 0.3 V superimposed with a sinusoidal perturbation voltage of 10 mV peak; ^gI_{MIN} is the minimum current achieved after 60 s to a single potential step of 1.2 V (i.e., the leak current); ^hThis potentiostat has a voltage range fixed at 10 V; ⁱPercentage error between the lowest and highest value.

Figure 3-3. Comparison of best and worst practices for the electrochemical characterisation of supercapacitors (Table 1, top), and reproducibility obtained using the best practices across different potentiostats (Table 2). Reproduced with permission from [359].

Although considered to be under steady state conditions, these tests remain sensitive to random errors arising from uncontrollable environmental variations. As previously introduced, changes in electrolyte composition during RDE cycling [357] can cause small deviations in the measured current or onset potential. Similarly, in GCD experiments, temperature fluctuations or electromagnetic noise can affect the charge-discharge profile of the materials being tested [359].

Sample-to-sample experiments introduce additional variability associated with inherent differences between samples. From the analysis presented above, testing multiple Pt/graphene electrodes prepared via the same protocol will inevitably reflect the influence of the handmade dropping affecting surface morphology, particle dispersion,

film thickness, and fabrication consistency on the resulting ORR activity. In the context of capacitance testing, different graphene electrodes, even if nominally identical, may exhibit variations due to differences in porosity, surface functionalisation, or differences in the mass of the electrodes [360].

Transient experiments investigate time-dependent processes and present specific challenges in capturing dynamic variable under controlled conditions. These techniques are particularly valuable in electrochemical systems, where reaction kinetics and mass transport evolve over time. Uncertainty in electrochemical transient experiments can arise from multiple sources. Cycling tests, which evaluate the lifetime performance of devices such as supercapacitors or catalysts, fall into this category. In this thesis, the focus will be on timewise and sample-to-sample experiments, leaving transient experiments as a topic for future work. Particularly relevant in this area is the work of Shangwei Zhou [361], who employed advanced diagnostics techniques, supported by Machine Learning, for fast and cost-effective online evaluation of the state of polymer fuel cells.

To systematically handle such diverse forms of uncertainty, it is useful to differentiate between repetition, which involves making repeated measurements under identical conditions, and replication, which refers to a structured approach to repetition that accounts for varying conditions. A hierarchical framework for replication in timewise experiments, originally proposed by Moffat [362], defines three key levels of replication: zeroth-order, first-order, and Nth-order replication. According to Coleman and Steele [352], this framework can also be applied also to sample-to-sample experiments.

At the zeroth-order level of replication, experimental conditions are held constant, and time is effectively treated as frozen. In practical terms, for electrochemical testing, this could include measuring properties such as electrical conductivity, electrode mass, or magnitudes at open-circuit potential. At this level, the only contributors to variability are the intrinsic noise and resolution of the instrumentation, which can be considered an irreducible source of error.

First-order replication introduces time as a variable while maintaining fixed instrumentation. In electrochemical testing, this level includes repeated measurements performed on the same sample or on different samples prepared using the same procedure (i.e. sample-to-sample experiments), which represents the standard practice for ensuring measurement reproducibility. When the same sample is used, this might involve measuring several polarisation curves on a single Pt/graphene electrode without re-polishing or re-preparing it, all within a narrow time interval to minimize environmental

drift. Similarly, performing several GCD cycles on the same graphene electrode under fixed current and electrolyte conditions would provide a first-order estimate of uncertainty.

When different samples are used, this includes acquiring polarisation curves for RDE tests or charge-discharge profiles in GCD tests from two samples prepared using the same procedure, measurement protocols, and data analysis techniques. In this context, variability in results can arise from the same factors previously described for sample-to-sample experiments, such as changes in ambient temperature or humidity, aging of the reference electrode, inconsistencies in catalyst ink deposition or electrode mixture, or the evolution of the electrode surface. Estimating uncertainty at the first-order level provides insight into the reproducibility of the measurement procedure and the reliability of the synthesis and sample preparation methods.

Nth-order replication introduces the highest degree of realism by allowing both instrumentation and experimental conditions to vary between measurements. Instrumentation refers to differences in the hardware used to perform the experiment, for example, different potentiostats, reference electrodes, reference electrode junctions, or rotation controllers. Experimental conditions, in contrast, refer to the aspects of the procedure and environment that affect the measurement but are not part of the instrument itself, such as electrolyte preparation, cell geometry, electrode mounting, temperature control, or laboratory-specific protocols.

A simple illustrative example is the reproduction of an RDE limiting-current measurement in two different laboratories: one lab may use a different potentiostat and a reference electrode from another supplier (differences in instrumentation), while preparing the electrolyte with different purity salts and using a slightly different electrode immersion depth (differences in experimental conditions). In this situation, systematic errors that would be normally fixed, such as voltage offset in a particular reference electrode or a bias arising from a specific electrode alignment procedure, become random variables when each measurement uses a different but nominally equivalent setup.

In RDE measurements, for instance, systematic deviation in measured limiting currents due to differences in electrolyte purity control or electrode geometry can be modelled statistically under Nth-order replication. Likewise, in GCD experiments, the capacitance calculated using different current sources, electrodes mounting systems or data analysis protocols [363] may exhibit dispersion that reflects broader systematic uncertainty. At this stage of replication, the uncertainty estimate integrates both random and systematic contributions, effectively representing the full range of expected experimental variation.

To practically apply the structured uncertainty framework discussed above, a series of RDE tests were conducted to investigate both the intrinsic performance of Pt/graphene-based catalysts and the variability introduced through sample preparation, testing conditions and laboratory equipment. These experiments offer a real-world case study of first-order and Nth-order replication, illustrating how differences in catalyst deposition, loading, and processing affect measured electrochemical performance.

3.2 Results

3.2.1 Experimental

Reduced Graphene Oxide (rGO) was synthesised using a modified Hummers method, as described in the Methods section, starting from Kish graphite from ArcelorMittal steelworks in Avilés, Spain. Reduction was carried out using ascorbic acid as the reducing agent. Elemental analysis revealed the following composition:

Table 3-1. Composition of Reduced Graphene Oxide from Kish graphite.

| | |
|------------------|---------------|
| %Carbon | 80-85% |
| %Oxygen | 15-20% |
| %Sulphur | 0.5-1% |
| %Nitrogen | <1% |
| %Ash | <1% |

This indicates that the degree of reduction was limited, and that the oxygen content remains relatively high in this sample. Figure 3-4 shows the morphology and Raman response of the reduced graphene oxide powder synthesised from Kish graphite. The SEM images reveal the characteristic wrinkled and crumpled-sheet structure of rGO, while the Raman spectrum displays the D and G bands typically observed in sp^2 carbon systems. The D band is associated with disorder and edge defects, and its relative intensity with respect to the G band provides information on the degree of structural imperfection. This figure confirms that the material obtained after chemical reduction remains highly defective. The pronounced D band and the resulting I_D/I_G ratio indicate the coexistence of multiple small sp^2 domains separated by oxygen-containing defects.

From a catalytic perspective, this structure is advantageous because the sp^2 network preserves electrical conductivity, while the defects and residual functional groups provide anchoring sites that promote strong metal-support interactions and the homogeneous dispersion of Pt nanoparticles. The SEM micrographs further show a crumpled, sheet like morphology consistent with a porous network of agglomerated graphene derived

layers. At both magnifications, the material appears to retain a degree of structural disorder and folding, which may help limit complete restacking of the sheets and preserve accessible surface area. Overall, such a morphology is generally favourable for use as a Pt catalyst support, as it can facilitate nanoparticle deposition and provide pathways for electronic conduction, although these observations should be regarded as qualitative.

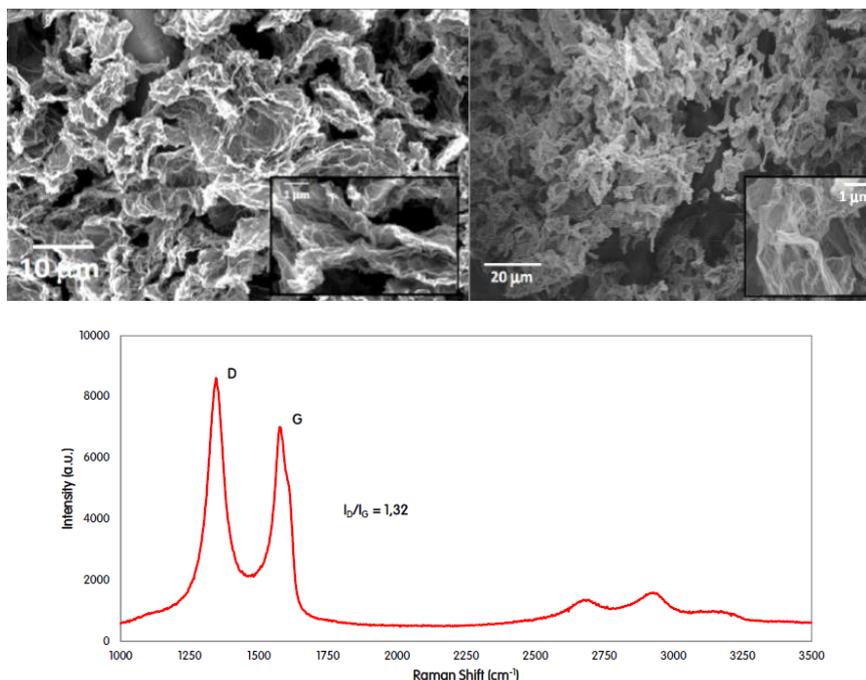


Figure 3-4. Top: SEM images of reduced graphene oxide powder; bottom: Raman spectra, with an I_D/I_G ratio typical from RGO powders, showing an intense D band related to a high presence of defects.

To deposit Pt on graphene, the polyol method described in Al-Youbi *et al.* [364] was followed. This method involves the reduction of a Pt-precursor in a liquid alcohol medium at moderate temperatures. The alcohol acts both reducing agent and solvent. First, a dispersion of reduced graphene oxide in ethylene glycol was ultrasonicated for 1h. Then, 1.25 mmol solution of H_2PtCl_6 in ethylene glycol was added dropwise under magnetic stirring at 60°C, to avoid the formation of Pt agglomerates. After 1 hour, a solution containing 12.5 mmol of $NaBH_4$ in ethylene glycol was added to reduce the Pt. The mixture was then refluxed at 130°C for 3 hours to ensure complete reduction. All steps were conducted under a nitrogen atmosphere and maintained at 50°C under flowing N_2 overnight. The graphene powder was then recovered by centrifugation, washed with ethanol and dried at 70°C for 8h.

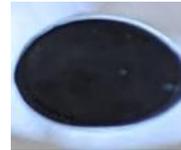
In addition to evaluating rGO as a support, expanded Kish graphite was also investigated as a catalyst. The goal was to explore the potential catalytic activity of iron-based

impurities present in the Kish. To increase its surface area, the expansion process described in Section 2.2.3 was carried out.

To evaluate the performance of these materials as catalysts for the oxygen reduction reaction (ORR), 6 mg of catalyst (Pt20%/rGO or Expanded Kish) was mixed with 200 μL of IPA, 800 μL of water, and 20 μL of Nafion. The ink was sonicated for 15 minutes in an ultrasonic bath. A 20 μL aliquot was then drop-cast onto a glassy carbon electrode (0.196 cm^2) in a rotating disk setup, resulting in a catalyst loading of 120 $\mu\text{gPt}/\text{cm}^2$ for the Pt20%/rGO. A similar ink was prepared using commercial Pt40%/Vulcan from Sigma-Aldrich. After deposition, the films were air dried. The visual appearance of the layers revealed inhomogeneous deposition (see Table 3-2), illustrating a key source of sample-to-sample variability characteristic of first-order replication.

To improve uniformity, different deposition and drying conditions were tested for both Pt20%/rGO and Pt40%/Vulcan samples. The results are summarized in Table 3-2.

Table 3-2. Deposition and drying conditions for different catalyst inks on top of working electrode prior to ORR activity evaluation.

| | Pt20%/RGO-1 | Pt40%/Vulcan-1 | Pt20%/RGO-2 | Pt40%/Vulcan-2 |
|------------------------|---|---|--|---|
| Concentration (mg/mL) | 5.88 | 0.59 | 2.94 | 2.94 |
| Water:IPA:Nafion ratio | 78:19:2 | 78:19:2 | 78:19:2 | 78:19:2 |
| Dropping | 20 μl | 20 μl | 5 μl x 4 times | 5 μl x 4 times |
| Drying | Air | Air | Infrared lamp | Ar flux |
| Visual aspect |  |  |  |  |

Prior to the electrochemical testing, the working electrode was cleaned by cycling the potential between 0 to 1.1 V vs. RHE at 50 mV/s for 20 cycles in N_2 -saturated 0.1 M HClO_4 electrolyte. The capacitive current was then measured at 10 mV/s. ORR polarisation curves were collected at 10 mV/s and 1600 rpm between 0.0 and 1.2 V vs. RHE in O_2 -saturated 0.1 M HClO_4 . The cathodic sweep from the CV was used for analysis, and the faradic current density was obtained by subtracting the capacitive current from the total measured current. Two samples of Pt20%/rGO and two of Pt40%/Vulcan were measured, while only one sample was tested using the expanded Kish ink. Inspired by the work of Tesch *et al.* [358], four key performance indicators (KPIs) were extracted from the curves to facilitate comparison:

- $E_0 = E(j = -0.1 \text{ mA} \cdot \text{cm}^{-2})$ [Onset potential]
- $E_1 = E(j = -1 \text{ mA} \cdot \text{cm}^{-2})$
- $E_3 = E(j = -3 \text{ mA} \cdot \text{cm}^{-2})$
- $E_5 = E(j = -5 \text{ mA} \cdot \text{cm}^{-2})$

Equation 3-1

The following metric was calculated as a measure of uncertainty for the first-order replicate pairs (Pt20%/rGO-1, Pt20%/rGO-2) and (Pt40%/Vulcan-1, Pt40%/Vulcan-2):

$$\text{Coef. Variation (\%)} = \left(\frac{\text{Standard Deviation}}{\text{Mean}} \right) \cdot 100 \quad \text{Equation 3-2}$$

Results are presented in Figure 3-5 and Table 3-3.

Table 3-3. Electrochemical parameters from polarisation curves for Expanded Kish Graphite, Pt20%/rGO and Pt40%/Vulcan samples.

| Sample | E_0 (V) | E_1 (V) | E_3 (V) | E_5 (V) |
|---|-----------|-----------|-----------|-----------|
| Expanded Kish Graphite | - | - | - | - |
| Pt20%/rGO-1 | 0.9998 | 0.8088 | 0.6107 | - |
| Pt20%/rGO-2 | 0.9377 | 0.8197 | 0.7187 | 0.5636 |
| <i>Coef. Variation^{rGO} (%)</i> | 5 | 1 | 11 | - |
| Pt40%/Vulcan-1 | 0.8597 | 0.7587 | 0.6367 | 0.1494 |
| Pt40%/Vulcan-2 | 0.9778 | 0.8117 | 0.6757 | 0.5457 |
| <i>Coef. Variation^{Vulcan} (%)</i> | 10 | 5 | 4 | 81 |

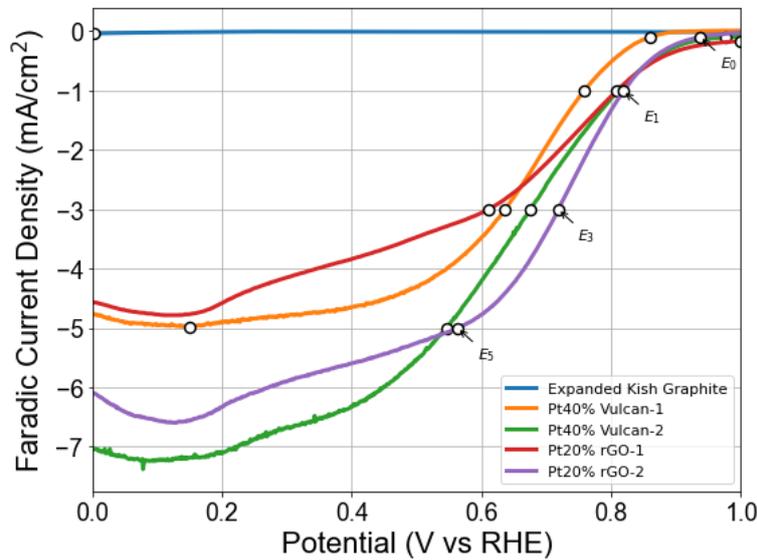


Figure 3-5. Polarisation curves for RGO20%Pt (red and purple lines), Vulcan40%Pt (orange and green lines) and Expanded Kish (blue line).

Figure 3-5 presents the ORR polarisation curves obtained for Pt supported on rGO, Pt supported on Vulcan, and expanded Kish graphite. The current density is plotted as a function of potential, enabling comparison of the onset of oxygen reduction and the current response at increasing overpotential.

This figure highlights the strong influence of both catalyst composition and film morphology on the ORR response. The similar onset potentials observed for the Pt-based catalysts suggest comparable intrinsic reaction kinetics at low current density. However, the curves diverge markedly at higher current densities, indicating increasing mass-transport limitations that depend strongly on film thickness, porosity and catalyst accessibility. As expected, expanded Kish graphite shows negligible ORR activity under acidic conditions.

Overall, the figure highlights that the preparation and drying of the catalyst ink play a decisive role in determining transport behaviour and measured performance.

Pt20%/rGO samples exhibited onset potentials of 0.9998 V and 0.9377 V, with a low Coefficient of Variation (CoV) of 5%, indicating consistent catalytic initiation behaviour. E_1 values also showed minimal variation (CoV = 1%), suggesting reproducible kinetics at low overpotentials. However, greater deviation at higher current densities (E_3 -CoV = 11%) likely reflects differences in film thickness and uniformity due to drying conditions. Films left to dry in ambient air undergo slow solvent evaporation, whereas infrared-lamp drying may produce faster, heat-assisted evaporation. Such differences could plausibly lead to variations in surface morphology or binder distribution, which in turn might influence the electrochemical response.

In contrast, Pt40%/Vulcan samples showed more variability. While onset potentials (0.8597 V and 0.9778 V) approached the value reported for Pt46%/Vulcan ($E_0=1.00$ V [365]), the E_5 potentials (0.1494 V and 0.5457 V) were much lower than the literature value (=0.8752 V), with an extremely high CoV of 81%. This suggests substantial mass transport or catalyst accessibility limitations, likely caused by inconsistent film morphology from differing drying methods (air vs. Ar flux).

These results highlight the critical role of ink formulation and drying strategy in achieving reproducible and high-performance ORR catalyst films, particularly at higher current densities where transport effects dominate.

The expanded Kish sample showed no measurable catalytic activity under the tested conditions, though it is possible that it could have some activity in alkaline conditions, where many Fe-based catalysts show much improved performance compared to acid electrolytes [366].

To evaluate the impact of catalyst loading and ink concentration on ORR performance, Pt40%/Vulcan samples were tested at a reduced loading of $30 \mu\text{gPt}\cdot\text{cm}^{-2}$, in contrast to the $120 \mu\text{gPt}\cdot\text{cm}^{-2}$ used in previous tests (Figure 3-5). Despite the lower loading,

polarisation curves in Figure 3-6 revealed that the limiting current density remained largely unaffected across most samples, indicating sufficient catalytic activity. However, the sample prepared with an ink concentration of 1.47 mg/mL and two drops exhibited much poorer performance, with a low onset potential ($E_0 = 0.7546$ V) and markedly depressed potentials at -1 mA·cm⁻² ($E_1 = 0.5117$ V) and -3 mA·cm⁻² ($E_3 = 0.1704$ V), likely due to non-uniform film formation or excessive thickness in localized regions.

Although the literature generally recommends Pt loadings in the range of 7-30 µg Pt·cm⁻², the tests presented in Figure 3-5 used a higher loading of 120 µgPt·cm⁻². Such high loading can produce a thick catalyst film, increasing mass-transport resistance. Thus, additional tests were performed on Pt40%/Vulcan at 30 µgPt·cm⁻², using different ink concentrations.

Figure 3-6 compares ORR polarisation curves for Pt/Vulcan electrodes prepared using different ink concentrations and deposition volumes at a lower Pt loading. The objective is to separate the effect of film morphology from that of metal mass. As shown in this figure, the lower loading had little effect on the limiting current density, except for the sample prepared with an ink concentration of 1.47 mg/mL and two drops, which exhibited reduced performance. Table 3-4 summarises the KPIs obtained.

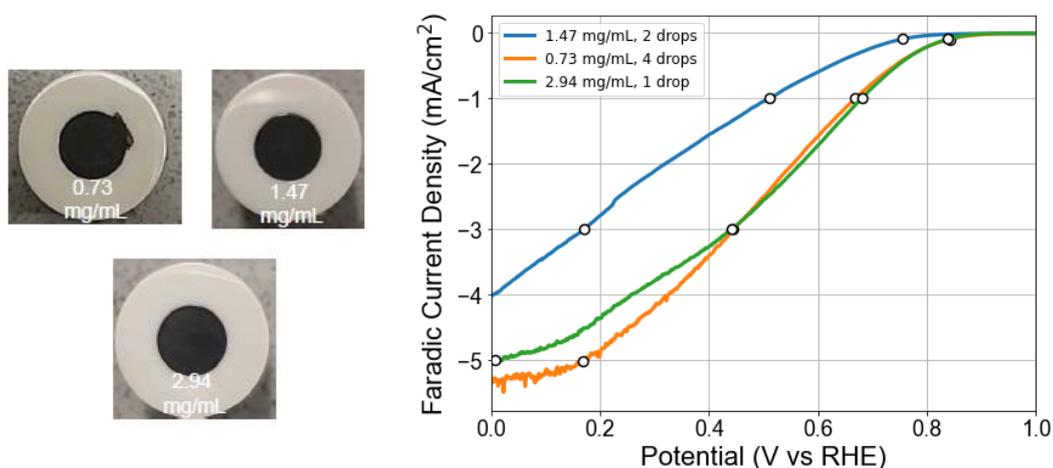


Figure 3-6. Left: photographs of catalyst films deposited on glassy carbon electrode under different preparation conditions. Right: Polarisation curves of Pt40%/Vulcan corresponding to a catalyst loading of 30 mg Pt·cm⁻² at different deposition conditions.

Table 3-4. Electrochemical parameters from polarisation curves for Pt40%/Vulcan replicates for different ink concentrations. *This Coef. Variation was only calculated for the pair (0.73 mg/mL, 4 drops) and (2.94 mg/mL, 1 drop).

| Sample | E_0 (V) | E_1 (V) | E_3 (V) | E_5 (V) |
|----------------------|-----------|-----------|-----------|------------|
| 1.47 mg/mL, 2 drops | 0.7546 | 0.5117 | 0.1704 | - |
| 0.73 mg/mL, 4 drops | 0.8416 | 0.6687 | 0.4447 | 0.1674 |
| 2.94 mg/mL, 1 drop | 0.8386 | 0.6807 | 0.4417 | 0.0074 |
| Coef. Variation (%) | 6 | 15 | 45 | - |
| Coef. Variation (%)* | - | - | - | 129 |

CoV analysis further highlights the disparity among samples. E_0 and E_1 showed moderate reproducibility (COV = 6% and 15%, respectively), while E_3 exhibited a high variation of 45%, emphasizing the sensitivity of higher current densities to film morphology. Notably, E_5 was only reached in two samples, with a remarkably high CoV of 129%, underscoring instability or mass-transport limitations at these current levels. These findings reinforce that, while reduced Pt loading can maintain ORR performance, careful optimisation of ink concentration and deposition volume is critical to ensure film uniformity and mitigate transport-related losses.

As discussed earlier, electrochemical parameters derived from RDE testing are highly sensitive to experimental conditions, including catalyst film morphology, loading, electrode coverage, electrolyte viscosity, and oxygen diffusion rate. Moreover, the skill of the experimentalist in preparing, depositing, and handling the samples introduces significant variability, emphasizing the practical relevance of the replication framework. To illustrate this effect, different deposition conditions were tested using Pt40%/Vulcan. The results, shown in Figure A1 (Appendix), highlight how even nominally identical materials can exhibit notable performance differences due to variations in preparation.

Finally, to explore the practical implications of Nth-order replication, ORR measurements were compared by conducting additional tests on identical Pt40%/Vulcan catalyst films in a second laboratory. The polarisation curves previously presented in Figure 3-5 were plotted alongside two new curves obtained from this independent setup.

Setup 2 differed from the original (Setup 1) in several key aspects. The ink formulation used a lower Nafion content (75% water, 25% isopropanol, and 0.02% Nafion), and film deposition was standardized to four 10 μ L drops from a 1 mg/mL ink suspension. Each drop was left to dry in air. Electrode conditioning was also more extensive, involving potential cycling under N_2 bubbling for 80 cycles at 70 mV/s, compared to 20 cycles at 50 mV/s in Setup 1. Moreover, a different potentiostat was employed (Gamry instead of BioLogic), introducing potential instrumental variability. Notably, rather than extracting the cathodic sweep from cyclic voltammetry, the ORR curve in Setup 2 was obtained using linear sweep voltammetry (LSV), offering a direct measurement of the reduction process. Figure 3-7 below shows photographs of the two instrumental setups used.

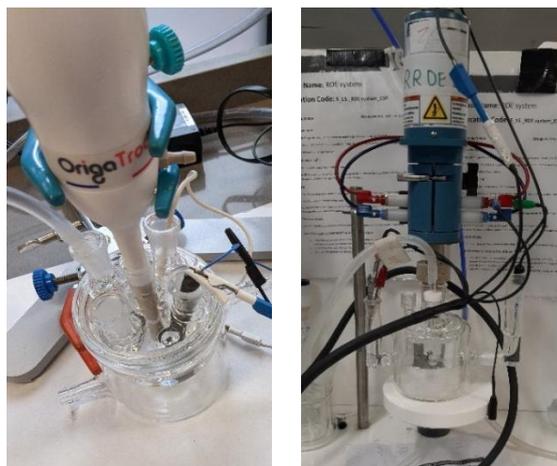


Figure 3-7. Experimental setups: Setup 1 (left) and Setup 2 (right).

Given these procedural and instrumental differences, the results from Setup 1 and Setup 2 cannot be interpreted as directly comparable in terms of absolute ORR performance. Instead, the two configurations represent distinct experimental conditions that introduce different sources of systematic variability. Within the framework of Coleman and Steele, Setup 2 does not constitute a zeroth-order or first-order replicate of Setup 1, but rather an Nth-order replicate, where multiple experimental parameters vary simultaneously.

The purpose of including Setup 2 in this chapter is therefore not to benchmark catalytic activity across laboratories, but to evaluate how performance dispersion evolves when higher-order variability is introduced. In this sense, the comparison between setups acts as a deliberate probe of uncertainty propagation, highlighting how changes in ink formulation, deposition, conditioning, and instrumentation manifest in the polarisation curves. The two Setup 1 curves shown previously in Figure 3-5 were obtained from films prepared with intentionally different ink concentrations (5.88 mg/mL and 0.59 mg/mL) to introduce controlled sample-to-sample variability. Under Setup 2, by contrast, both films were prepared using the same ink formulation and deposition protocol and therefore constitute replicates under tightly controlled conditions. It is important to emphasize that the Setup 1 and Setup 2 measurements are not replicates of one another; rather, they reflect Nth-order variability arising from the combined differences in ink formulation, deposition procedure, conditioning protocol, and instrumentation. The four resulting polarisation curves (two from Setup 1 and two from Setup 2) are presented together in Figure 3-8, with the corresponding quantitative metrics summarised in Figure 3-8 demonstrates that changing the experimental setup leads to measurable shifts in onset potential and current density. These shifts do not necessarily reflect intrinsic catalyst behaviour, but rather differences in ink formulation, film deposition, conditioning and instrumentation. The figure therefore provides practical evidence that ORR performance

metrics must be interpreted within an uncertainty framework, particularly when comparing results across laboratories.

Table 3-5.

Figure 3-8 demonstrates that changing the experimental setup leads to measurable shifts in onset potential and current density. These shifts do not necessarily reflect intrinsic catalyst behaviour, but rather differences in ink formulation, film deposition, conditioning and instrumentation. The figure therefore provides practical evidence that ORR performance metrics must be interpreted within an uncertainty framework, particularly when comparing results across laboratories.

Table 3-5. Summary of KPIs extracted from ORR polarisation curves of Pt40%/Vulcan catalyst films measured under two experimental setups.

| Sample | Setup | E_0 (V) | E_1 (V) | E_3 (V) | E_5 (V) |
|----------------------------|-------|-----------|-----------|-----------|-----------|
| Pt40%/Vulcan 1 | 1 | 0.8597 | 0.7587 | 0.6367 | 0.1494 |
| Pt40%/Vulcan 2 | 1 | 0.9778 | 0.8117 | 0.6757 | 0.5457 |
| Pt40%/Vulcan 3-1 | 2 | 0.7872 | 0.7012 | 0.6552 | 0.6032 |
| Pt40%/Vulcan 3-2 | 2 | 0.7592 | 0.6932 | 0.6472 | 0.5872 |
| <i>Coef. Variation (%)</i> | | 12 | 7 | 3 | 46 |

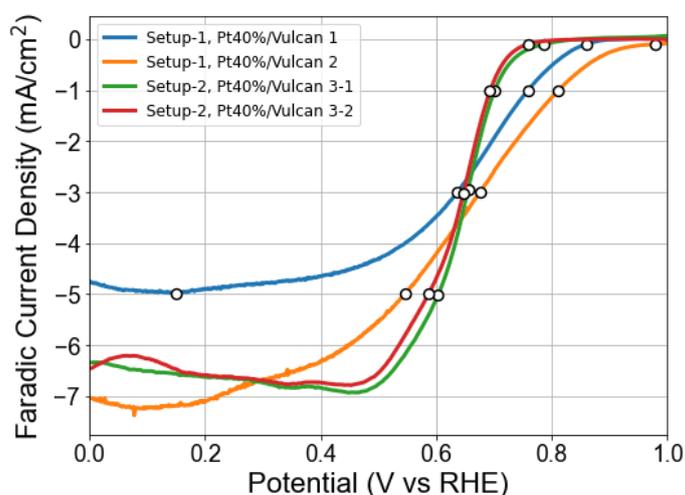


Figure 3-8. ORR polarisation curves for Pt40%/Vulcan catalyst films measured under two experimental setups.

The onset potentials (E_0) ranged from 0.759 V to 0.9778 V, with a CoV of 12%, reflecting moderate variability in kinetic onset across conditions. As the current density increased, the results became more consistent. At -1 and -3 mA·cm⁻², the CoVs decreased to 7% and 3%, respectively, indicating a convergence of behaviour in the kinetic-to-diffusion transition region. This suggests that the intrinsic catalytic activity of the Pt40%/Vulcan material remains comparable under both setups despite procedural differences. At -5 mA·cm⁻², however, variability increased substantially (CoV = 46%). This large

spread likely stems from a combination of mass-transport limitations and differences in film thickness but may also reflect variations in instrumental setup (e.g., potentiostat type), chemical materials (e.g., electrolyte purity, IPA and Nafion concentration), and electrochemical technique (cyclic voltammetry vs. linear sweep voltammetry) between the two laboratories. A detailed analysis of the individual sources of variability is beyond the scope of the present work.

These findings reinforce the importance of structured replication in electrochemical metrology. The observed variability across samples and tests aligns with the first- and Nth-order replication levels, providing a concrete demonstration of how procedural differences and operator-dependent factors influence experimental outcomes. Careful consideration of these uncertainties is essential to ensure robust and reproducible assessments of electrocatalyst performance.

This concludes the ORR studies using the RDE technique. Due to the high variability observed in this approach and the limited ORR activity of Kish graphite, the work that follows focuses on reduced graphene oxide capacitance evaluation, which offers improved reproducibility for subsequent analyses, as will be shown in the following section.

3.2.2 Galvanostatic Charge-Discharge Tests

As described in the literature review, the reduction of graphene is a key step in producing an effective supercapacitor electrode for several reasons. On the one hand, the removal of oxygen-containing groups within the crystalline domains of graphene can result in the formation of holes in the graphene flakes. This structural modification ultimately affects the porosity of the electrode and its specific surface area. On the other hand, the degree of reduction determines the carbon-to-oxygen ratio, which in turn influences the electrical conductivity, pseudocapacitance and double-layer capacitance of the material.

Commercial graphene oxide from Sixth Element was used as the raw material. The chosen reduction method, chemical reduction with ascorbic acid, offers a scalable, non-toxic, and environmentally friendly route suitable for producing graphene-based supercapacitors. To investigate measurement uncertainty and reproducibility, rGO electrodes were synthesized under three distinct thermal conditions: 65 °C for 3 hours, 80°C for 30 minutes and 95 °C for 30 minutes. In addition, to introduce a controlled source of experimental noise representative of typical laboratory variability, two common stirring methods were compared: magnetic stirring and mechanical paddle stirring at 350 rpm. In line with prior descriptions of experimental uncertainty, this can be classified as a sample-to-sample variation corresponding to first-order replication.

GCD measurements were used to assess the reproducibility of specific capacitance under these conditions at three current densities: 1 A/g, 2 A/g and 10 A/g. Four key KPIs were extracted from the discharge profiles, corresponding to the times at which the voltage reached 0 V, -0.4 V, -0.8 V, and -1.2 V vs. a silver reference electrode:

- $t_0 = t(E = 0 V)$
 - $t_{0.4} = t(E = -0.4 V)$
 - $t_{0.8} = t(E = -0.8 V)$
 - $t_{1.2} = t(E = -1.2 V)$
- Equation 3-3

For clarity, the variables used in the uncertainty analysis can be expressed in terms of input factors x and measured responses y . For each GCD experiment, the set of input variables is defined as:

$$\mathbf{x} = [T_{red}, t_{red}, \text{stirring method}, j] \quad \text{Equation 3-4}$$

where

- T_{red} : reduction temperature $\in \{65 \text{ }^\circ\text{C}, 80 \text{ }^\circ\text{C}, 95 \text{ }^\circ\text{C}\}$
- t_{red} : reduction time $\in \{3 \text{ h}, 0.5 \text{ h}\}$
- Stirring method $\in \{\text{magnetic bar}, \text{paddle at } 350 \text{ rpm}\}$
- j : applied current density $\in \{1, 2, 10 \text{ A g}^{-1}\}$

The measured outputs constitute the response vector:

$$\mathbf{y} = [t_0, t_{0.4}, t_{0.8}, t_{1.2}, C, \text{CoV}] \quad \text{Equation 3-5}$$

where $t_0, t_{0.4}, t_{0.8}$ and $t_{1.2}$ are the characteristic discharge times defined in Equation 3-3; C is the specific capacitance obtained from the discharge slope; and CoV is the coefficient of variation used to quantify reproducibility (Equation 3-2).

It is worth noting that this explicit x and y formulation was not applied in the earlier RDE section, as that study focused primarily on qualitative uncertainty sources rather than on structured input-output modelling. Here, however, such notation is introduced because the GCD results form the experimental dataset used in the ML optimisation presented in the next chapter.

Each reduced graphene oxide sample was mixed with PTFE and conductive carbon black C65™ (Timcal) in a weight ratio of 80:10:10 to form the positive electrode in a three-electrode configuration. As counter-electrode, activated carbon Norit™ (Cabot Corporation) was used, also in an 80:10:10 composition with PTFE and C65™, with a 10:3 mass ratio relative to the rGO material. A 1M KOH aqueous solution served as the electrolyte, and a silver bar was used as reference electrode. All measurements were

conducted in a Swagelok-type electrochemical cell. Figure 3-9 shows the Swagelok-type three-electrode cell used for galvanostatic charge-discharge measurements of the rGO electrodes. The working, counter and reference electrodes are assembled inside the cell together with the alkaline electrolyte.

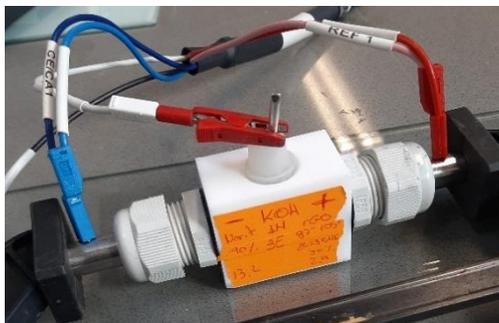


Figure 3-9. Swagelok cell used to extract galvanostatic charge-discharge (GCD) curves with a three-electrode setup.

Figure 3-10 presents the GCD curves comparing magnetic stirring and paddle stirring under various conditions. Each row corresponds to a different reduction condition: 65 °C for 3 h (top row), 80 °C for 0.5 h (middle row), and 95 °C for 0.5 h (bottom row). Each column represents a different current density: 1 A/g (left), 2 A/g (middle), and 10 A/g (right). Within each plot, two GCD curves are shown: one corresponds to the use of a magnetic stir bar during the reduction step, while the other corresponds to the use of a stir paddle rotating at 350 rpm. White circles mark the time points corresponding to specific potentials as defined in Equation 3-3. As in the case of the RDE experiments, the coefficient of variation (Equation 3-2) was calculated to quantify the uncertainty associated with these tests.

Results for 2 A/g are presented in Table 3-6, while those for 1 A/g and 10 A/g are included in Table A1 in Appendix A. Across all conditions, the 2 A/g tests showed the lowest variation, with CoV values ranging from 0.5% to 7%. However, the effect of current density on reproducibility was not uniform across reduction conditions. At 65 °C for 3 h, 1 A/g exhibited the highest CoVs (up to 14%, see Table A1 in Appendix A), suggesting greater sensitivity to the stirring method under slower reduction kinetics associated with lower-temperature, longer-duration treatment. In contrast, under the 80 °C for 0.5 h condition, the CoVs for 2 A/g (up to 7%) were higher than those for 1 A/g. At 95 °C for 0.5 h, 10 A/g showed the greatest variability, with CoVs reaching up to 9%. These results indicate that no single current density consistently exhibited the highest or lowest variability, and that reproducibility is influenced by a combination of current density, reduction conditions (temperature and time), and stirring dynamics.

In contrast, RDE measurements conducted under controlled sources of variability exhibited substantially higher and more erratic CoV values, ranging from 1% to as high as 129%. This high dispersion, particularly at voltages associated with higher current densities, suggests that data-driven optimisation using RDE may require substantially more experimental repetitions to achieve reliable results.

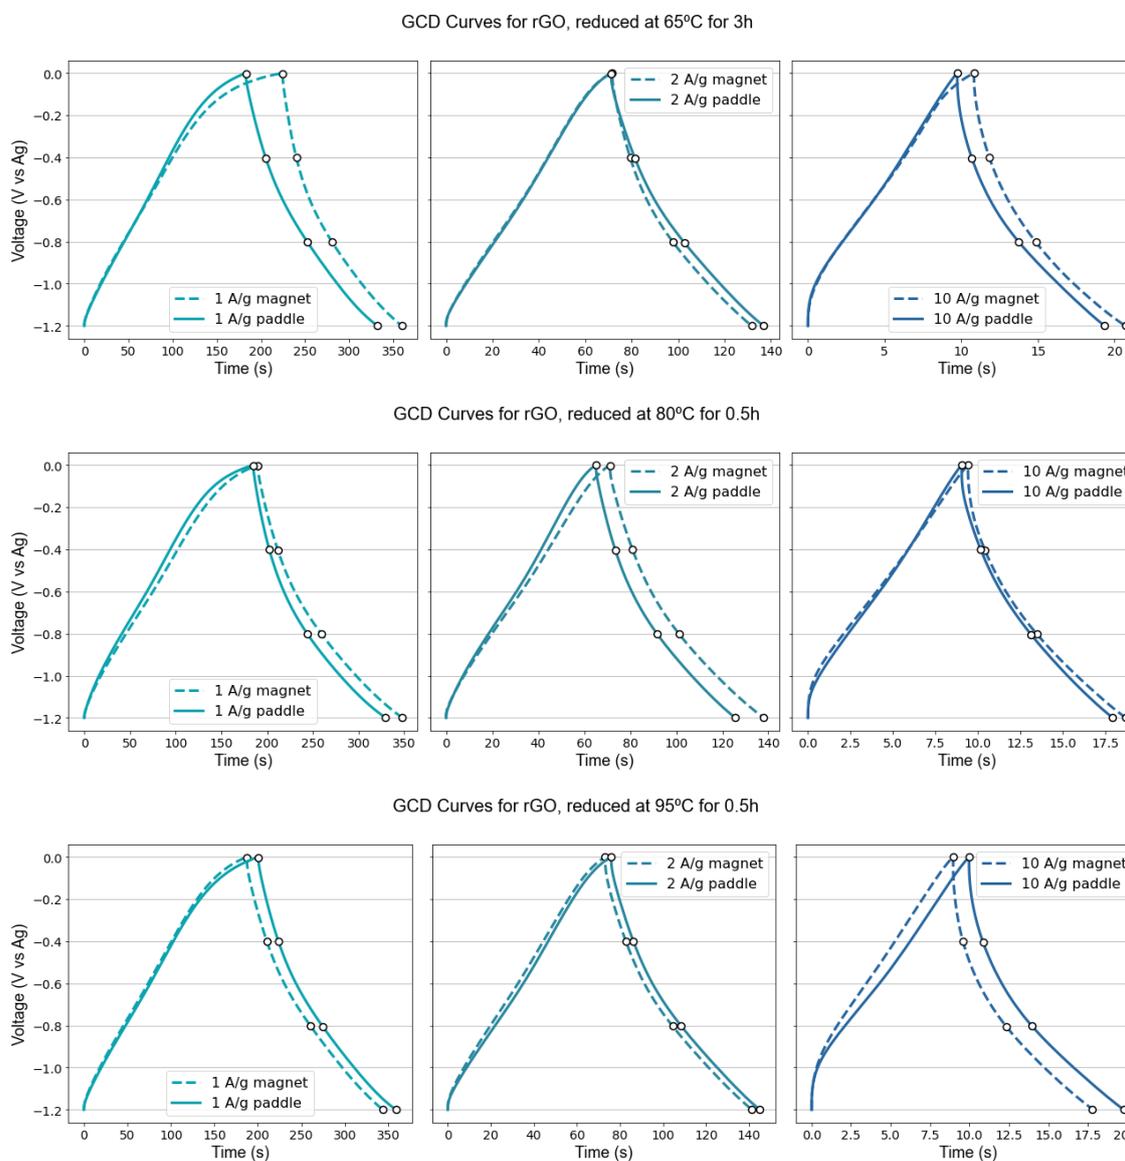


Figure 3-10. GCD curves under different reducing conditions and mechanical stirring method, measured at current densities of 1 A/g (left), 2 A/g (centre), and 10 A/g (right).

Table 3-6. Discharge times (in seconds) at different applied voltages (0, -0.4, -0.8, -1.2 V) for rGO-based electrodes tested at 2 A/g, fabricated under three reduction conditions (65 °C, 3 h; 80 °C, 0.5 h; 95 °C, 0.5 h) using two stirring methods (magnetic and paddle). The coefficient of variation (CV, %) between

stirring methods is shown for each condition, illustrating the reproducibility of electrode performance across processing variables.

| Reduction conditions | Stirring | t_0 (s) | $t_{0.4}$ (s) | $t_{0.8}$ (s) | $t_{1.2}$ (s) |
|----------------------------|----------|--------------|---------------|---------------|---------------|
| 65°C, 3h | Magnet | 71.63 | 79.36 | 97.84 | 131.71 |
| 65°C, 3h | Paddle | 71.13 | 81.34 | 102.7 | 136.87 |
| <i>Coef. Variation (%)</i> | | <1 | 2 | 3 | 3 |
| 80°C, 0.5h | Magnet | 71.18 | 80.99 | 101.33 | 137.93 |
| 80°C, 0.5h | Paddle | 65.09 | 73.54 | 91.67 | 125.46 |
| <i>Coef. Variation (%)</i> | | 6 | 7 | 7 | 7 |
| 95°C, 0.5h | Magnet | 73.13 | 82.81 | 104.51 | 141.3 |
| 95°C, 0.5h | Paddle | 75.94 | 86.25 | 108.57 | 145.14 |
| <i>Coef. Variation (%)</i> | | 3 | 3 | 3 | 2 |

Thus, given the overall low variation and reproducibility of GCD tests, particularly at 2 A/g, this technique was selected for use in Gaussian Process modelling and Bayesian Optimisation experiments, which are described in the following chapter.

Finally, specific capacitance values were calculated from the GCD curves using Equation 3-6,

$$C = I \frac{\Delta t}{\Delta V \cdot m} \quad \text{Equation 3-6}$$

where Δt is the discharge time extracted from the GCD curves and $\Delta V = 1.2V - IR \text{ Drop}$. The IR drop has been systematically extracted for each curve using a custom computational routine implemented in Python [367]. Figure 3-11 and Table 3-7 summarises the obtained results.

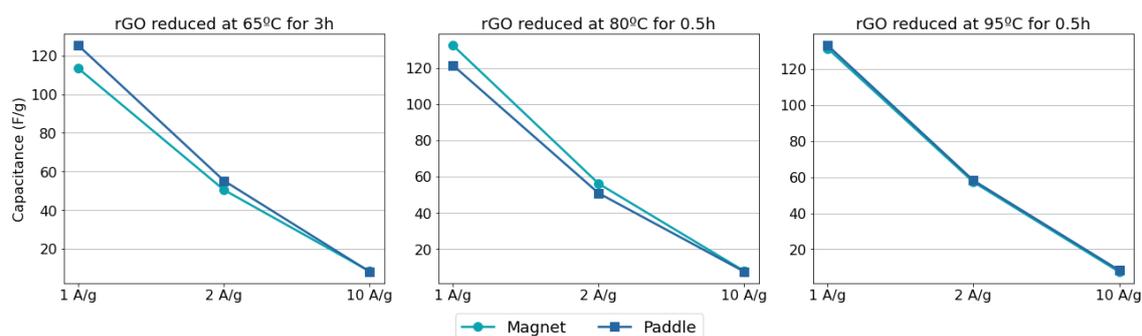


Figure 3-11. Specific capacitance values calculated using Equation 3-6 for all tested conditions.

Table 3-7. Summary of specific capacitance values calculated and CoVs for each condition.

| Reduction conditions | Stirring | <i>C</i> (F/g) @ 1A/g | <i>C</i> (F/g) @ 2A/g | <i>C</i> (F/g) @ 10A/g |
|-----------------------------|-----------------|------------------------------|------------------------------|-------------------------------|
| 65°C, 3h | Magnet | 113.35 | 50.51 | 8.44 |
| 65°C, 3h | Paddle | 125.25 | 55.26 | 8.19 |
| <i>Coef. Variation (%)</i> | | 7 | 6 | 2 |
| 80°C, 0.5h | Magnet | 132.33 | 56.10 | 7.84 |
| 80°C, 0.5h | Paddle | 121.35 | 50.83 | 7.54 |
| <i>Coef. Variation (%)</i> | | 6 | 7 | 3 |
| 95°C, 0.5h | Magnet | 131.40 | 57.32 | 7.53 |
| 95°C, 0.5h | Paddle | 133.09 | 58.22 | 8.39 |
| <i>Coef. Variation (%)</i> | | 1 | 1 | 8 |

The values of specific capacitance followed the expected trend of decreasing with increasing current density across all conditions. At 1 A/g, capacitance values were the highest, with values typically exceeding 120 F/g, whereas at 10 A/g, the capacitance dropped sharply, often falling below 10 F/g. This trend is characteristic of porous carbon-based materials, where limited ion access at high discharge rates reduces the effective use of surface area.

The analysis of the stirring method as a source of variability revealed relatively modest differences in capacitance between magnetically stirred and paddle-stirred samples. For example, under the 80 °C, 0.5 h condition, capacitance at 1 A/g was 132.33 F/g for the magnet-stirred sample and 121.35 F/g for the paddle-stirred one, yielding a CoV of 6%. At 95 °C, the values were 131.40 F/g and 133.09 F/g, respectively, with a remarkably low CoV of 1%, indicating high reproducibility. For the 65 °C condition, the capacitance values at 1 A/g were slightly lower overall but followed a similar pattern, with the paddle-stirred sample yielding a higher value (125.25 F/g) than the magnetically stirred one (113.35 F/g), corresponding to a CoV of 7%.

These differences, while present, remained within a relatively narrow range and did not suggest any systematic or disruptive effect of the stirring method on the resulting electrochemical properties. The highest value observed (133.09 F/g) is superior to that of a typical commercial activated carbon reference (86 F/g at 0.5 A/g [244]), although it remains well below the best-performing materials reported within the family of graphene-based electrodes (up to 429.7 F/g [215]).

At intermediate (2 A/g) and high (10 A/g) current densities, the variations between the two stirring approaches were similarly minor. Capacitance values at 2 A/g ranged from ~50 to 58 F/g with CoVs between 1% and 7%, while values at 10 A/g remained consistently low (7.5–8.5 F/g), with CoVs between 2% and 8%. In all cases, the overall pattern of lower capacitance with increasing current density was preserved, and no

evidence was observed that either stirring method introduced relevant anomalies in performance.

Qualitative inspection of the GCD curves indicated that the majority of the measurements exhibited highly symmetric charge–discharge profiles, consistent with efficient and reversible capacitive behaviour across the dataset. Curves at 1 A/g showed the least symmetry, likely due to the pseudocapacitive contribution of rGO at lower current densities, which tends to introduce kinetic asymmetries and reduce Coulombic efficiency. No consistent difference in profile shape or charge retention behaviour was observed between stirring methods, further supporting the notion that the induced noise does not introduce strong electrochemical artefacts.

Importantly, while some variation in capacitance values is present across different reduction conditions, this study does not attempt to isolate or explain the effects of temperature and time on rGO performance. The dataset remains too limited to confidently attribute these effects to the reduction conditions, as they may also reflect random experimental variation. As such, differences between samples reduced at 65°C, 80°C, or 95°C are interpreted here only within the narrow context of replication noise and not as evidence of fundamental trends. A more detailed investigation of reduction conditions and their influence on structure and electrochemical behaviour, interlaboratory studies, will be undertaken in the next chapter of this thesis.

3.3 Conclusions and Future Work

Electrochemical characterisation is essential for assessing materials in energy conversion and storage, but meaningful interpretation requires a clear understanding of both measurement and experimental uncertainty. These uncertainties arise not only from instrumentation limits but also from factors such as sample preparation, environmental conditions, and testing protocols. In this chapter, these challenges are explored through two experimental directions: (1) oxygen reduction reaction (ORR) studies of raw Kish and Kish-derived graphene, and (2) capacitance evaluation of reduced graphene oxide. Both efforts aim to assess performance metrics while examining the impact of reproducibility and variability, informed by literature on catalyst benchmarking and data-driven material optimisation.

Inspired by the works of Moffat and Coleman & Steele on uncertainty analysis, a series of RDE and GCD tests were conducted in which controlled sources of variability were intentionally introduced. The goal was to evaluate the sensitivity of these electrochemical techniques to different forms of uncertainty and to assess the reproducibility of performance data under realistic conditions. To achieve this, the coefficient of variation

was employed as a metric to quantify uncertainty across timewise, sample-to-sample experiments, including first-order and N-order replicates.

RDE measurements revealed substantial variability across samples, with CoV values ranging from as low as 1% to as high as 129%. This high dispersion, especially at higher current densities, underscores the sensitivity of ORR performance to subtle differences in film morphology, catalyst loading, ink composition, and deposition conditions. Even under controlled procedures, reproducibility remained limited, particularly in the mass-transport region. These findings suggest that data-driven optimisation using RDE methods is subject to considerable uncertainty and may require extensive replication to produce reliable and comparable results across different setups and laboratories.

GCD measurements appear to offer a robust and reproducible approach for evaluating the capacitive behaviour of rGO materials under controlled variability. The impact of the chosen noise source, namely the stirring method, was observable but modest, with CoV values generally ranging from 1% to 7%, and reaching a maximum of 14% under specific conditions. This suggests that GCD can reliably tolerate typical sample-to-sample variability without introducing large or unpredictable shifts in output. This contrasts with previous findings using RDE methods, where small changes in ink deposition or drying conditions led to much larger deviations in the polarisation curves. Given the overall low variation observed in GCD, this technique was selected for the optimisation experiments presented in the next chapter.

While the next chapter will examine the influence of reduction temperature and time on the structural and electrochemical properties of rGO using GCD, the RDE technique will not be further considered in this thesis. Nevertheless, the high variability observed in RDE results underscores two key directions for future research: quantitative film characterisation, using techniques such as SEM, AFM, or profilometry, to correlate morphology with performance, and the standardisation of deposition protocols. Implementing automated or robotic dispensing could help reduce operator-induced variability and enhance the reproducibility of RDE measurements across different laboratories.

4 Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes

Summary

In this chapter, a Gaussian Process (GP) is used as a surrogate model to apply Bayesian Optimisation (BO) for identifying the temperature and time reduction conditions that maximise capacitance of reduced graphene oxide (rGO). An experimentally determined White Noise was incorporated into the GP as a fixed measurement-noise term, approximating it to a grey-box problem. This noise was quantified from controlled experiments targeting two dominant sources of variability, namely heating instability during reduction and electrode-mass differences. Although these sources do not represent the full range of experimental variability, they capture the main contributors to the broader uncertainty introduced across the experimental chain, including the reduction, electrode preparation and electrochemical characterisation steps. The GP obtained after an initial set of tests is evaluated using Leave-One-Out Cross-Validation, demonstrating the ability to predict each experimental capacitance within a 95% confidence interval. Furthermore, the pseudocapacitance behaviour of the rGO electrodes is evaluated using Nyquist plots and cyclic voltammetry. Finally, a second set of tests is incorporated into an updated version of the GP, which is adapted to account for the fact that the new data points correspond to measurements performed under different setup conditions. Thanks to the inclusion of an additional term in the exponential component of the kernel, the accuracy of the GP remains comparable to that obtained with the initial set of tests.

The innovativeness of this work stems from three key factors. First, the relationship between the time and temperature of the chemical reduction and the capacitance of reduced graphene oxide, using a non-toxic, environmentally friendly reducing agent such as ascorbic acid, has not yet been investigated. Second, this is the first time in which a GP has been modulated by experimentally determined noise in the context of graphene-based supercapacitors. Finally, physics-informed bounds are imposed on the length-scales associated with synthesis conditions and measurement set-ups during model fitting, ensuring alignment with the constraints imposed by scaling up the reduction process and the generalisability of the results.

4.1 Introduction

The process of reducing graphene oxide is well known and different routes have been deeply studied since the early days of graphene research [322,368–371]. These include

thermal treatments [372–375], microwave radiation [318,376,377] or even the use of microorganisms [378,379]. One of the most cited and easier to scale is the chemical route [380–383]. Numerous chemicals have been described as effective reducing agents for graphene oxide, including nitrogen-containing reducing agents such as hydrazine, borohydrides such as NaBH_4 or oxygen-based reductants such as methanol or ascorbic acid. Although the efficiency of many of these chemicals for reducing graphene oxide has been evaluated extensively, the influence of the reaction conditions such as temperature and time with respect to the reagent used is less known.

Cheng and Yan [384] studied the reduction of graphene oxide using N,N-dimethylacetamide (DMAc) dispersed in water at atmospheric pressure. They evaluated the FT-IR spectra of the produced rGO after four different combinations of temperature and time of reduction; the best reduction was produced when the reaction was carried out at 150°C for 5 hours. The study was limited to four conditions defined by three temperatures (100°C, 120°C and 150°C) and only two different reaction times (1h and 5h). On the other hand, the evaluation of the reduction efficiency was supported by the IR spectra peaks, so neither percentage of carbon and oxygen in the samples, nor other features such as specific surface area were reported. In addition, the reagent employed is toxic and considered a restricted substance in the European market [385], which limits the scalability of the results.

Wan *et al.* explored the influence of reducing agent, reaction time and temperature of rGO-based aerogels for oil sorption [386]. Although for the influence of the reducing agent, the final chemical composition of the rGO was evaluated with elemental analysis, FT-IR spectroscopy, X-ray diffraction and Raman spectroscopy, the part of the study focused on evaluating the influence of reaction time and temperature was limited to SEM micrographs and to oil sorption capacity assessment.

Guex *et al.* studied how different reduction times and temperatures affected the electrical conductivity of reduced graphene oxide, using sodium borohydride [319]. The work was extensive, covering up to nine reducing times from 2 min to 24 h, coupled with four reduction temperatures: 20°C, 40°C, 60°C and 80°C. The influence of the drying step was also addressed by comparing vacuum filtration, evaporation and freeze drying. This showed that the quality of the film formation is key to obtain good electrical conductivity, and best conductive films were obtained with vacuum filtration. Even though the temperature and time windows chosen were compatible with a scalable process, the reagent, sodium borohydride, is a toxic and flammable solvent, which presents a limitation when planning to reproduce these results at larger scales.

When comparing their results with other works, using the same reducing agent and similar reduction temperatures and times, the authors highlighted a discrepancy between the carbon / oxygen ratios and the conductivity of the films, which did not correlate positively, as one may expect. They speculated that this might come from the sensitivity of the conductivity to the preparation of the films and the measuring techniques, together with the potential for unaccounted contamination or humidity.

Comparable results are essential to consistently explain how synthesis conditions influence the properties of the final materials. This is particularly important for nanomaterials, where target properties are highly sensitive to unquantified noise sources. As described in Chapter 3, such sources include structural and morphological variations caused by human manipulation of samples, environmental effects affecting the material's chemistry, and inconsistencies in measurement techniques and data analysis across laboratories. These challenges are part of the broader reproducibility crisis in science, where difficulties in replicating experimental findings hinder progress and innovation across multiple disciplines [387–389].

Evaluating electrochemical properties, such as capacitance, is especially vulnerable to inconsistencies arising from variations in measurement protocols and procedures [332,363,390–395]. Moreover, the synthesis conditions at larger scales are typically more constrained than those at laboratory scales for several reasons. These include the need to guarantee operational safety, costs limitations associated with the capital expenditure required for scaling up, and even the lack of knowledge in chemical engineering for some of the innovative synthesis procedures developed in the laboratory [396]. All the above challenges pose a major obstacle to scaling up high-value materials, which are critical for the transition to a zero-emissions economy.

In this regard, the increasing adoption of Machine Learning (ML) techniques to guide materials research can be explained by how such techniques support the treatment of experimental noisy data and physics-constrained processes [397–401]. In particular, kernel-based methods such as Gaussian Processes, offer a practical method to project to and from high-dimensional space. Unlike other kernel methods like Support Vector Machines, which do not provide uncertainty quantification, Gaussian Processes do, as they are inherently probabilistic models. This enables the use of Bayesian Optimisation to assist experimentalists in the Design of the Experiments, handle noisy data (aleatoric uncertainty), quantify what remains unknown (epistemic uncertainty) and manage experimental or simulated stochastic processes [402–406].

As detailed in Section 2.3.2, the hyperparameters of a GP can be optimised to fit the available data as closely as possible, using techniques such as cross-validation or MLE. On the other hand, the Squared Exponential Kernel (Equation 2-9) provides high flexibility with a simple form defined by just two hyperparameters. However, this approach is entirely data-driven and treats the relationship between input variables and the function as a black-box, where no prior domain knowledge is incorporated.

In the present work, prior knowledge is introduced through the experimentally determined uncertainty associated with capacitance measurements. For this reason, the noise variance was not treated as a hyperparameter during optimisation but fixed to the experimentally obtained value and added to the diagonal of the covariance matrix. By constraining this component of the GP using experimentally measured information rather than inferring it solely from the data, the resulting model can be regarded as a grey-box GP. Based on this model, Bayesian Optimisation (BO) is subsequently applied to identify the optimal synthesis parameters for maximising capacitance. The surrogate model is later updated to incorporate the uncertainty arising from Nth-order replication, as formally defined in Section 3.1, and later applied to the experimental context of the present study in Section 0.

In this work, replication captures multiple sources of experimental uncertainty typical of graphene-based electrode fabrication and characterisation, such as variations in temperature control during reduction, electrode mass, and measurement-related noise. While the magnitude of these uncertainties may be non-negligible, particularly when data originate from different experimental set-ups or laboratories, variability on the order of 10-20% between nominally identical replicates is considered acceptable in this work, based on prior experimental experience with similar materials and characterisation protocols. Rather than invalidating the use of GP models, such variability is explicitly incorporated as an observation-noise term, allowing the GP to propagate uncertainty into both predictions and acquisition decisions. This probabilistic treatment enables BO to remain robust to moderate experimental noise, favouring parameter regions that are not only high-performing but also reliable under realistic experimental conditions.

The results are structured into three sections. In the first section, the uncertainties associated with synthesis parameters, specifically temperature and mass fluctuations, are quantified and incorporated into a constrained modelling framework based on physical considerations, ensuring process scalability.

The second section introduces a GP model built from an initial set of experiments. The model is evaluated through Leave-One-Out Cross-Validation (LOOCV) and refined via

BO over successive iterations, enhancing the predictive mapping between synthesis conditions and electrode capacitance. The section concludes with an electrochemical characterisation of the rGO electrodes, using Nyquist plots and cyclic voltammetry to assess their pseudocapacitance behaviour.

Finally, the third section addresses interlaboratory variability by estimating a new noise term corresponding to measurements from a different setup and a new term in the GP kernel which enables the adaption of the GP to systematic offsets.

4.2 Results

4.2.1 Formal Definition of the Gaussian Process Regression Model

In this study, Gaussian Process (GP) regression is employed to model the relationship between graphene oxide reduction conditions and the experimentally measured specific capacitance. The regression problem is defined as follows.

Each experiment is associated with an input vector

$$\mathbf{x}_i = [T_i, t_i]^T \in \mathbb{R}^2 \quad \text{Equation 4-1}$$

where $i = 1, \dots, n$ indexes the individual experimental observations, T_i denotes the reduction temperature ($^{\circ}\text{C}$) and t_i denotes the reduction time (h). The corresponding output $y_i \in \mathbb{R}$ represents the experimentally measured specific capacitance obtained from galvanostatic charge-discharge measurements at a current density of 2 A/g. The dataset is denoted by

$$\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \quad \text{Equation 4-2}$$

with the input matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ formed by stacking the input vectors row-wise and the output vector $\mathbf{y} \in \mathbb{R}^n$ containing the corresponding capacitance measurements. In this study, the input dimensionality is $d = 2$.

The measured output is assumed to be a noisy observation of an underlying latent function $f(\mathbf{x})$, such that

$$y_i = f(\mathbf{x}_i) + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_n^2), \quad \text{Equation 4-3}$$

where ε_i represents additive observation noise with variance σ_n^2 . In this work, σ_n^2 is not treated as a free model parameter but is estimated empirically from experimental variability arising from temperature control and electrode mass fluctuations, as described below in section 4.2.2.

A Gaussian Process prior is placed over the latent function,

$$f(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}')), \quad \text{Equation 4-4}$$

where $k(\mathbf{x}, \mathbf{x}')$ is the anisotropic squared-exponential kernel defined in Equation 2-7, here applied to the two-dimensional input vector $\mathbf{x} = [T, t]^\top$, allowing independent length-scales for temperature and reduction time.

In several cases, multiple experiments are performed under nominally identical synthesis conditions, resulting in repeated input vectors $\mathbf{x}_i = \mathbf{x}_j$ with different measured outputs $y_i \neq y_j$. Within the GP framework, such replicated observations are naturally accommodated through the additive observation-noise term.

The covariance matrix of the observed data is given by

$$\mathbf{K}_y = \mathbf{K}_f + \sigma_n^2 \mathbf{I}, \quad \text{Equation 4-5}$$

where \mathbf{K}_f is the noise-free kernel matrix evaluated using $k(\mathbf{x}, \mathbf{x}')$ at the inputs \mathbf{X} , and \mathbf{I} is the identity matrix. The presence of $\sigma_n^2 \mathbf{I}$ ensures that \mathbf{K}_y remains positive definite even when multiple observations share identical inputs, preventing numerical singularities and allowing replicated measurements to be interpreted as independent noisy samples of the same latent function value.

Consequently, replicated experiments reduce the posterior uncertainty at the corresponding input locations while the posterior mean converges toward the noise-weighted average of the observed outputs. This behaviour is consistent with the experimental design adopted in this chapter, where replicated reductions using different heating configurations are intentionally introduced to quantify temperature-induced variability and propagate it into the GP model.

4.2.2 Uncertainty Quantification and Model Constraints

An analysis of the potential sources of experimental noise arising from the graphene oxide reduction and electrode preparation steps is described below. This analysis is limited to two different reducing temperature points (80°C and 120°C) and a single reducing time point (0.5 hours). Rather than covering the entire temperature and time ranges considered for the capacitance model, the study is focused on a local case and, in the absence of sufficient replicates across all operating conditions, assumes homoscedasticity. In other words, the measurement noise is treated as approximately uniform across the explored synthesis conditions, recognising that potential variations in error at different temperature-time combinations cannot be fully resolved with the

available data. The experimental determination of uncertainty described here is based on a selection of potential sources of noise associated with lab tasks and procedures that are typical for this type of material and its characterisation. Thus, the values obtained may be applicable to other research groups working on capacitance measurements of materials from the same family, enhancing the extrapolation of the results.

To clarify the origin and propagation of the experimental uncertainties considered in this study, Figure 4-1 presents a schematic overview of the potential sources of error underlying the capacitance measurements. The scheme distinguishes between uncertainties arising during electrode preparation, graphene oxide reduction stage (temperature control) and electrochemical characterisation. These contributions are progressively quantified throughout the chapter; uncertainties related to temperature and mass are evaluated first, while inter-laboratory variability is introduced and quantified later in Section 4.2.4.

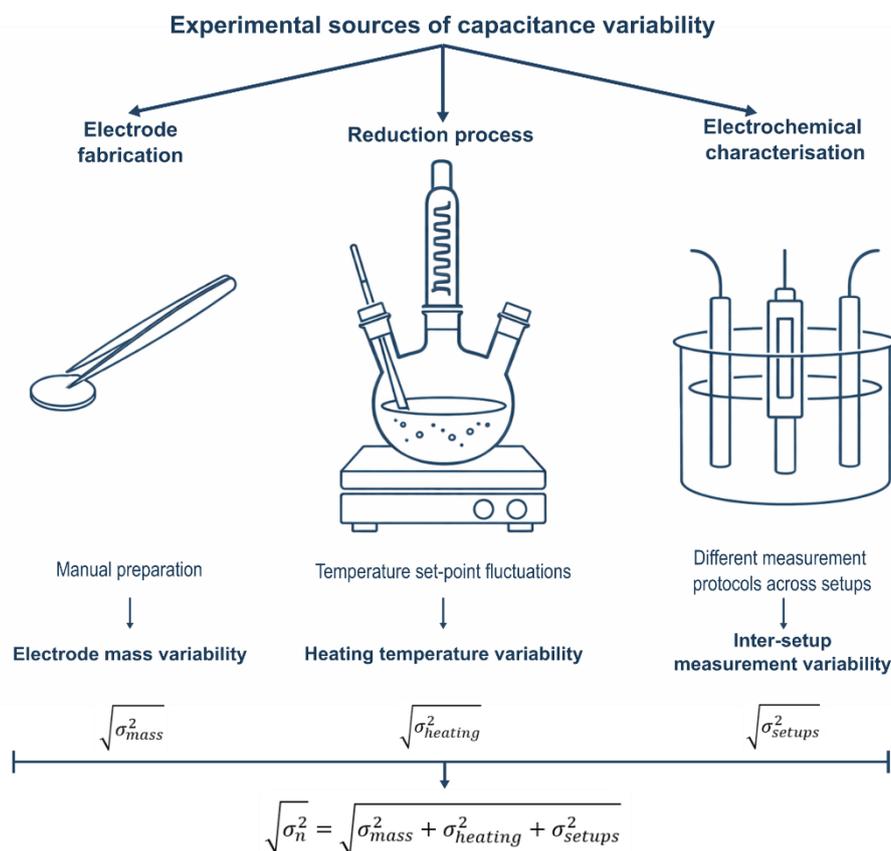


Figure 4-1 Conceptual scheme illustrating the main sources of experimental uncertainty affecting the measured specific capacitance. The contributions considered include mass variability during electrode preparation, temperature fluctuations during graphene oxide reduction and measurement-protocol variability.

In a reduction process based on the chemical route at fixed temperature, a typical lab procedure consists of introducing the reactants in a flask placed on a hot plate, as

illustrated in left photo of Figure 4-2. The increase in the temperature within the flask is limited by connecting it to a condenser, whose cold reservoir is an ice bath. This helps maintain a constant water volume during the reduction.

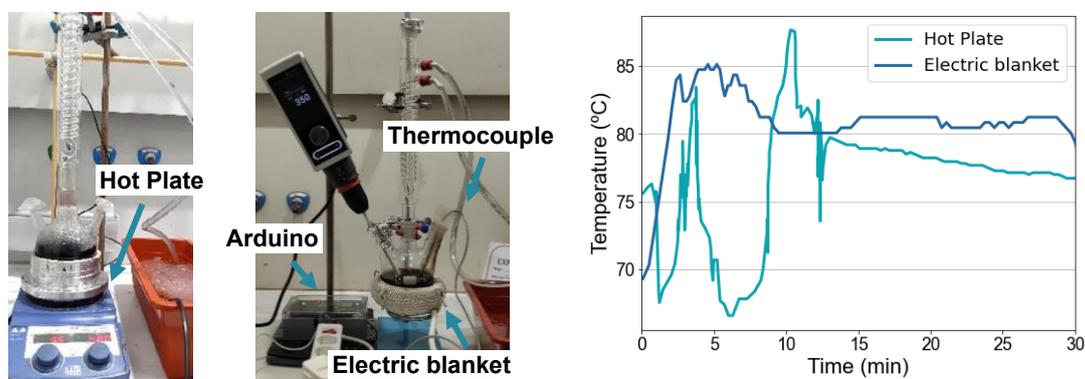


Figure 4-2. Experimental set-ups for the reduction of graphene oxide using ascorbic acid as reducing agent under controlled temperature. Left: set-up with temperature control via a hot plate. Middle: setup with temperature control via an Arduino-based system, which measures bath temperature using a thermocouple and adjusts power to an electric blanket accordingly. Right: temperature and time curves recorded for both set-ups during an identical reduction treatment at 80°C for 30 mins.

The heating power released by the hotplate is controlled by an internal setpoint, influenced by the temperature achieved by the plate. However, this may be a source of error, as the objective is to develop a model where the input is the temperature of the reduction bath on the capacitance, and not the temperature of the hotplate used for the reduction. To assess this, a thermocouple was introduced inside the bath during a reduction at 80°C for 30 min. As can be seen from the right plot of Figure 4-2, there is a sharp, irregular fluctuation in the temperature of the bath inside a window between -10°C and +5°C from the goal temperature.

To try to estimate the uncertainty this may cause in the capacitance, a new set-up is configured (Figure 4-2, middle). In this set-up, the heating power is provided by an electric blanket, in contact with the three-neck flask. The blanket is connected to an Arduino, which modulates the released power according to the set temperature measured by a thermocouple installed inside the bath. Thanks to this, a smoothening of the temperature profile is achieved, as can be seen from the right picture of Figure 4-2.

To assign an uncertainty in the capacitance caused by the temperature fluctuations, two reductions were carried out at 80°C for 30 mins, one where the temperature is controlled by the hotplate and another at the same conditions, but where the temperature is controlled by the system based on the electric blanket. This approach is similar to that described in Chapter 3, where the differences in the outputs between replicates were used to determine the coefficient of variation. Specifically, in Chapter 3 the coefficient of variation was computed by performing a series of first-order replicates under deliberately noisy conditions and quantifying the relative spread of the measured responses. This

procedure isolates the variability that arises from the measurement system and from controlled perturbations introduced to expose experimental sensitivity. As in that case, a source of uncertainty is deliberately introduced through first-order replicates performed under nominally identical experimental conditions in order to estimate experimental noise. This noise is incorporated into the grey-box Gaussian process model as an additive observation-noise term, such that the measured capacitance can be written as

$$y = f(\mathbf{x}) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma_n^2), \quad \text{Equation 4-6}$$

where y denotes the experimentally measured specific capacitance corresponding to the synthesis conditions \mathbf{x} , $f(\mathbf{x})$ represents the latent (noise-free) capacitance response of the system, and σ_n^2 is the observation-noise variance estimated from the variability observed between replicates. Accordingly, in this study the output of the Gaussian Process corresponds to the experimentally measured specific capacitance y , while the input variables \mathbf{x} , describe the synthesis conditions. By applying the same strategy in the present chapter, with two different sources of heat for the reductions serving as intentional sources of variability, an empirical measure of the temperature-induced variability in capacitance is obtained and propagated as σ_n^2 in the grey-box GP model.

The electrodes were prepared as follows: first, the components were weighed and mixed on a glass plate. Then, the wet mixture was placed between two pieces of paper, pressed against a glass, and rolled slightly to produce a slab with a homogeneous thickness. Using a circular cutter, the slab was then cut into individual circular electrodes. Finally, the electrodes were placed on a glass plate covered with holed aluminium foil and left to dry overnight at 80°C and 150 mbar. Figure 2-8 illustrates the entire process.

To compare the capacitances of rGO-containing electrodes from different heating set-ups, electrodes with the most similar masses were characterised to minimise the impact caused by mass differences on capacitance measurements. Thus, an electrode weighing 2.9 mg with rGO reduced using a hot plate was compared to an electrode weighing 2.998 mg with rGO reduced using an Arduino-controlled electric blanket. The results are shown in Figure 4-3 and Table 4-1.

Figure 4-3 compares the galvanostatic charge–discharge (GCD) curves of electrodes fabricated using rGO reduced either on a conventional hot plate or using an Arduino-controlled electric blanket. At 1 and 2 A·g⁻¹, both electrodes display the expected triangular GCD profile, indicating capacitive behaviour. However, the electrode prepared with the hot plate shows longer discharge times at all current densities. This suggests

that the electrode produced using the blanket experiences stronger rate limitations, possibly due to differences in porosity, contact resistance or microstructure induced by the heating protocol.

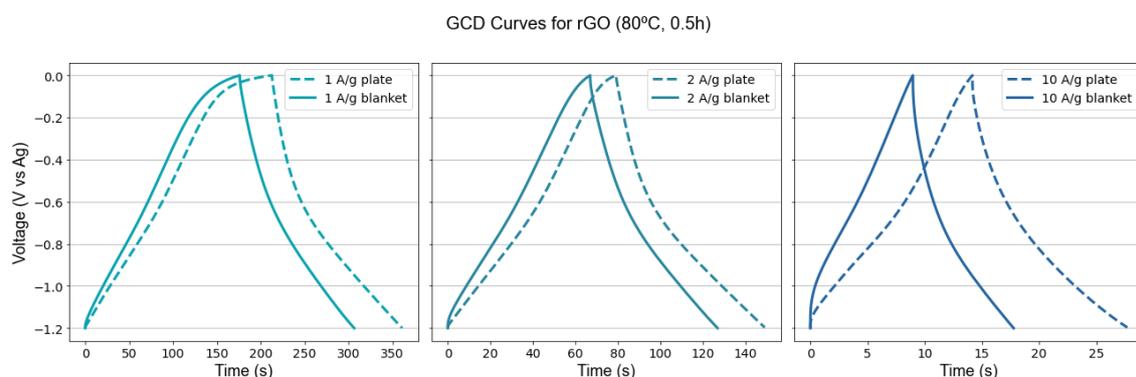


Figure 4-3. GCD curves of rGO electrodes reduced using two different heating systems. The curves compare electrodes whose graphene oxide was reduced at 80°C for 30 mins using either a conventional hotplate (dashed line) or an Arduino-controlled heating blanket (solid lines). Galvanostatic charge-discharge tests were performed at 1 A/g (left), 2 A/g (middle) and 10 A/g (right).

Table 4-1. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for electrodes with rGO from reductions using two different heating procedures. The bottom row corresponds to the standard deviation between them.

| Current Density Heating system | 1 A/g | 2 A/g | 10 A/g |
|-----------------------------------|-------------------|-------------------|-------------------|
| Electric blanket | 109.69 F/g | 100.77 F/g | 74.84 F/g |
| Hot plate | 124.74 F/g | 117.55 F/g | 115.28 F/g |
| $\sqrt{\sigma_{heating}^2}$ | 10.64 F/g | 11.87 F/g | 28.60 F/g |

The second source of uncertainty considered in this study is the fluctuation in mass of the prepared electrodes. Given the difficulty of fabricating electrodes with identical masses, due to the influence of manual labour in the process, this represents a plausible source of discrepancies between experiments. Consequently, mass fluctuations are a reasonable contributor to the underlying variance in a model that aims to relate the temperature and duration of graphene oxide reduction to electrode capacitance.

To do so, two different electrodes were prepared, both using reduced graphene oxide (rGO) as the active material, reduced at 80°C for 30 mins. The quantities of rGO used were slightly different: one electrode contained 2.75 mg of rGO, while the other contained 2.46 mg. In both cases, rGO made up 80% of the electrode mass, with the remaining 20% consisting of 10% carbon black (Timcal Super C65) as conductive additive and 10% PTFE as a binder. The results are shown in Figure 4-4 and

Table 4-2.

The curves are very similar in shape at all current densities, indicating that the electrodes behave comparably despite the small mass difference. This is reflected quantitatively in

Table 4-2: the specific capacitances differ only marginally at 1 and 2 A·g⁻¹ and remain within a few percent of one another even at 10 A·g⁻¹. The low standard deviations (≤ 7 F·g⁻¹) suggest that, once the current is normalised to mass, moderate differences in active-material loading introduce only limited dispersion into the calculated capacitance. Nevertheless, the slightly larger deviation at 10 A·g⁻¹ again points to high-rate measurements being more sensitive to small geometric or structural variations between electrodes.

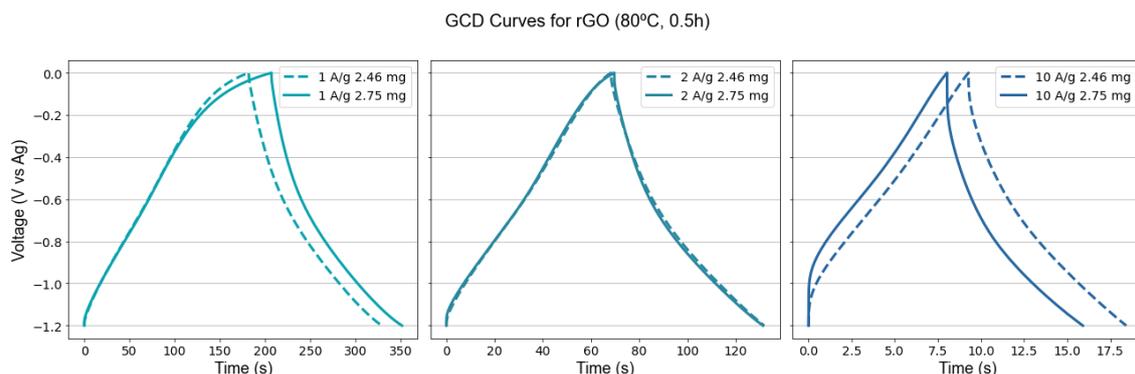


Figure 4-4. GCD curves for reduced graphene oxide electrodes. The graphene oxide was reduced at 80°C for 30 mins. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). The dashed lines correspond to electrodes containing **2.46 mg** of rGO, while the solid lines are electrodes with **2.75 mg** of rGO.

Table 4-2. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes with different mass. The bottom row corresponds to the standard deviation between them.

| Current Density Electrode mass | 1 A/g | 2 A/g | 10 A/g |
|-----------------------------------|-------------------|-------------------|------------------|
| 2.46 mg | 123.65 F/g | 106.79 F/g | 77.88 F/g |
| 2.75 mg | 121.39 F/g | 103.92 F/g | 68.08 F/g |
| $\sqrt{\sigma_{mass}^2}$ | 1.60 F/g | 2.03 F/g | 6.93 F/g |

Together, these results indicate that processing-related variables such as heating history and electrode handling can measurably influence the apparent capacitance, particularly at high current density. In contrast, moderate fluctuations in electrode mass contribute much less to the overall uncertainty once the data are expressed in specific terms. These observations support the decision to treat fabrication-related variability explicitly in the uncertainty framework developed in this work.

Assuming these sources of uncertainty are reasonably independent, justified by the fact that they arise from distinct equipment and procedures, namely temperature control during reduction and mass measurement during electrode preparation, the total uncertainty is calculated as the square root of the sum of the squared individual uncertainties:

$$\sqrt{\sigma_n^2} = \sqrt{\sigma_{heating}^2 + \sigma_{mass}^2} \quad \text{Equation 4-7}$$

Table 4-3. Calculation of the total noise for 1 A/g, 2 A/g and 10 A/g.

| | 1 A/g | 2 A/g | 10 A/g |
|---------------------|------------------|------------------|------------------|
| $\sqrt{\sigma_n^2}$ | 10.76 F/g | 12.04 F/g | 29.42 F/g |

Each uncertainty level was estimated from the standard deviation of two curves (corresponding to two electrodes), which may not fully capture the underlying noise distribution due to the limited sample size. This is acknowledged as a limitation, although the GP framework can still operate with these uncertainty estimates as inputs.

Because these uncertainties were derived empirically, contributions from inherent measurement variability (i.e., irreducible error) are assumed to be embedded within the reported values. The irreducible error was estimated experimentally by repeating galvanostatic charge-discharge measurements on two nominally identical electrodes prepared under the same synthesis and testing conditions (identical rGO reduction temperature and duration, heating setup, and closely matched electrode masses). Since temperature and mass variations were minimised by design, the remaining dispersion in the measured specific capacitance was attributed to irreducible measurement variability. The irreducible error was therefore quantified as the standard deviation between the capacitance values obtained for the two electrodes at each current density. Full experimental details and results are provided in Appendix B (page 199).

Notably, at 2 A/g and 10 A/g, the irreducible error is smaller than the uncertainty levels associated with temperature and mass fluctuations ($\sigma_{temperature}, \sigma_{mass}$), supporting the assumption that these sources dominate the overall uncertainty under these conditions. However, at 1 A/g, the irreducible error slightly exceeds σ_{mass} , suggesting that under certain conditions, measurement variability may play a greater role in the overall uncertainty, without compromising the validity of the data, as this variability is explicitly quantified and incorporated into the uncertainty analysis. Nevertheless, since all uncertainty estimates are derived from empirical trials, the irreducible error is considered to be adequately accounted for in the reported values. Finally, although uncertainty was calculated for three different current densities, the GPs presented in this study are restricted to data obtained at 2 A/g. This choice simplifies the analysis and keeps the modelling within the scope of the thesis, while allowing a clear demonstration of the

methodology. The investigation of uncertainty dependence and GP modelling across different current densities is left for future work.

As stated in Section 2.3.2, the library used in this thesis is Scikit-learn, a popular open-source machine learning library in Python. In Scikit-learn, the Gaussian Processes are fit to the data following *Algorithm 2.1*. as presented in [348]. In brief, it consists of calculating the probability of the observed points given a model with a set of hyperparameter values (also known as marginal likelihood or evidence) and find the combination of hyperparameters that maximise such probability.

In this study, \mathbf{x}_i represents the vector of experimental input variables, namely the reduction temperature and reduction time, and y_i represents the measured specific capacitance for that input. Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ denote the matrix of input points. The noisy covariance matrix is constructed as

$$\mathbf{K}_y = \sigma_f^2 \exp\left(-\frac{1}{2} \mathbf{Q}\right) + \sigma_n^2 \mathbf{I} \quad \text{Equation 4-8}$$

where \mathbf{Q} is a matrix with entries $Q_{ij} = (\mathbf{x}_i - \mathbf{x}_j)^\top \mathbf{M}(\mathbf{x}_i - \mathbf{x}_j)$, and the exponential function is applied element-wise. Here, $\mathbf{M} = \text{diag}(l_1^{-2}, \dots, l_d^{-2})$ is a diagonal metric matrix encoding the inverse squared length-scales for each input dimension, derived from the length-scale vector $\boldsymbol{\ell} = (l_1, \dots, l_d)$. The marginal likelihood in log form is:

$$\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}) = -\frac{1}{2} \mathbf{y}^\top \mathbf{K}_y^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K}_y| - \frac{n}{2} \log 2\pi \quad \text{Equation 4-9}$$

where $\boldsymbol{\theta} = \{\sigma_f^2, \boldsymbol{\ell}, \sigma_n^2\}$ represents the kernel hyperparameters. In this thesis, the GP is implemented using the Scikit-learn library, as previously stated. The observation-noise variance σ_n^2 is fixed to the empirically estimated experimental uncertainty while σ_f^2 and the length-scale vector $\boldsymbol{\ell}$ are varied during the optimisation process, with a range of bounds by default that spans from 10^{-5} until 10^5 . In this thesis, the noise term σ_n used in the GP is taken directly from the experimentally derived uncertainty (Equation 4-7); consequently, σ_n is not independently estimated by the model but is computed from the empirical variability observed in the heating and mass measurements.

An effective use of GP requires the selection of the right kernel, that introduces the prior knowledge from the scientific domain. To follow this philosophy, and given that our aim is to build a model representing the capacitance response to variations in graphene synthesis conditions, Equation 4-7 is computed as σ_n^2 in the GP. Similarly, the bounds of

l are limited to capture a more representative model. To do so, the following heuristic rule is considered [407]:

$$l \approx \frac{\sqrt{\sigma_{X_j}^2}}{\lambda}, \quad \lambda \in [1, 10], \quad j = 1, \dots, d \quad \text{Equation 4-10}$$

where l denotes the characteristic length-scale associated with the j -th input variable, $\sigma_{X_j}^2$ is the empirical variance of the input variable X , and λ is a dimensionless scaling factor controlling the smoothness of the latent function. Equation 4-10 restricts the value of the length-scale to a value that is shorter than the ‘spread’ of the input variables, with the aim of capturing the information about the objective function ϕ that can be obtained across the input space. For the temperature input, $\sqrt{\sigma_{temperature}^2}$, is computed over a restricted operating window (65-95 °C), corresponding to the temperature range suitable for scalable graphene production in a pilot-plant environment (see Figure 4-5 below).



Figure 4-5. Graphene production pilot plant. Photo courtesy of TheNextPangea®.

Additionally, a time window from 0.5h until 3h is applied to ensure a cost-effective reduction process, with the experimental space filled by data points varying at 5°C and 30-minute intervals, respectively. This gives the following values for $\sqrt{\sigma_X^2}$:

$$\sqrt{\sigma_{temperature}^2} = 10.80 \Rightarrow l_{temperature} \in [1.08, 10.80] \quad \text{Equation 4-11}$$

$$\sqrt{\sigma_{time}^2} = 0.94 \Rightarrow l_{time} \in [0.09, 0.94] \quad \text{Equation 4-12}$$

4.2.3 Surrogate Modelling and Experimental Optimisation

In addition to the capacitance given by 80°C and 0.5 h, two more pairwise points of reduction time and temperature were evaluated, given by the following conditions:

- 65°C, 3 hours
- 95°C, 0.5 hours

The GCD curves and capacitances are summarized in Figure 4-6 and Table 4-4, respectively. With this selection of hyperparameters and initial points, a GP model was constructed. The results are shown in Figure 4-7. The top plot represents the predicted mean distribution, with crosses indicating the observation points. As expected, the mean remains close to zero across most of the plotted region, except in areas near the observed data points. These localized variations in the mean are influenced by the length-scales of the covariance function Equation 4-11 and Equation 4-12. The anisotropic shape of the GP’s response around the observed data points, which exhibits a more extended spread along the Time axis, reflects the shorter length-scale associated with this variable Equation 4-12.

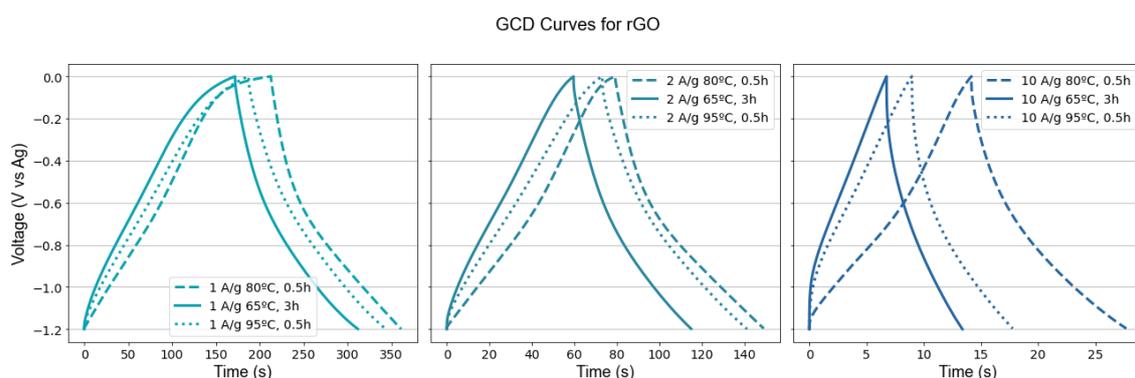


Figure 4-6. GCD curves for reduced graphene oxide electrodes under three different reduction conditions. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). The dashed lines correspond to 80°C and 0.5 hours, the solid line to 65°C and 3 hours, and the dotted line to 95°C and 0.5 hours.

Table 4-4. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes produced under three different reduction conditions: 80°C for 5 hours, 65°C for 2 hours and 85°C for 1.75 hours.

| Current Density Reduction T ^a , time | 1 A/g | 2 A/g | 10 A/g |
|--|-------------------|-------------------|-------------------|
| 80°C, 0.5h | 124.74 F/g | 117.55 F/g | 115.28 F/g |
| 65°C, 3h | 117.81 F/g | 104.55 F/g | 78.22 F/g |
| 95°C, 0.5h | 131.40 F/g | 114.64 F/g | 75.28 F/g |

At first glance, the GP’s mean prediction of zero capacitance in unobserved regions may seem physically unrealistic, as capacitance is generally expected to have nonzero values. However, a GP does not provide only a point prediction but also models the confidence associated with that prediction. In regions far from the observed data, the predictive uncertainty becomes large, indicating that the model assigns low confidence to the mean value.

Figure 4-7 shows the Gaussian-process surrogate model built to describe the dependence of capacitance on reduction temperature and time. The upper plots display the predicted mean capacitance over the design space, while the lower plots show the associated predictive uncertainty.

The bottom plot in Figure 4-7 visualises the 95% confidence interval, which extends by nearly 200 F/g in regions of high uncertainty. This indicates that while the GP mean might suggest an unphysical prediction (zero capacitance), the high uncertainty corrects this by acknowledging the model's lack of knowledge in those regions. In other words, while the GP predicts zero in those regions, such a prediction is not a confident estimate, reflecting the fact that the model has not learned enough from the available data points to make a reliable prediction.

With only three experimental observations, the Gaussian Process (GP) model is not intended to provide a physically accurate global prediction of capacitance. Instead, its purpose at this stage is to demonstrate how limited experimental data are incorporated into a probabilistic surrogate model that simultaneously represents predictions and associated uncertainty across the design space.

The figures demonstrate that the GP not only interpolates between measured points but also quantifies the level of confidence in each prediction. Regions far from experimental data exhibit higher uncertainty, while densely sampled regions show narrower confidence intervals. This behaviour is essential for guiding Bayesian optimisation and for ensuring that model-based decisions account for experimental noise.

To apply Bayesian Optimisation, the Expected Improvement (EI) acquisition function was selected, as it effectively balances exploration and exploitation during the optimisation process [408]. Figure 4-8 presents the EI values for temperature and time, indicating that the next optimal point for graphene oxide reduction is around 80°C and just above 0.5 hours, as highlighted by the red star in the graph. However, it is important to note that this prediction is based on a limited number of data points and is intended to illustrate the application of the method at this stage rather than to provide a definitive prediction.

This figure is particularly important because it links modelling with experimental decision-making. Rather than sampling the design space uniformly, EI directs new experiments toward regions where the balance between high predicted capacitance and high uncertainty is most favourable.

The optimisation process was carried out until an experimental budget of ten additional points was reached. The corresponding results are presented in Figure 4-9. Measurements were repeated four times at 80°C, 0.5 h and twice at 72°C, 0.8 h, resulting in a total of 13 observations across 9 unique points.

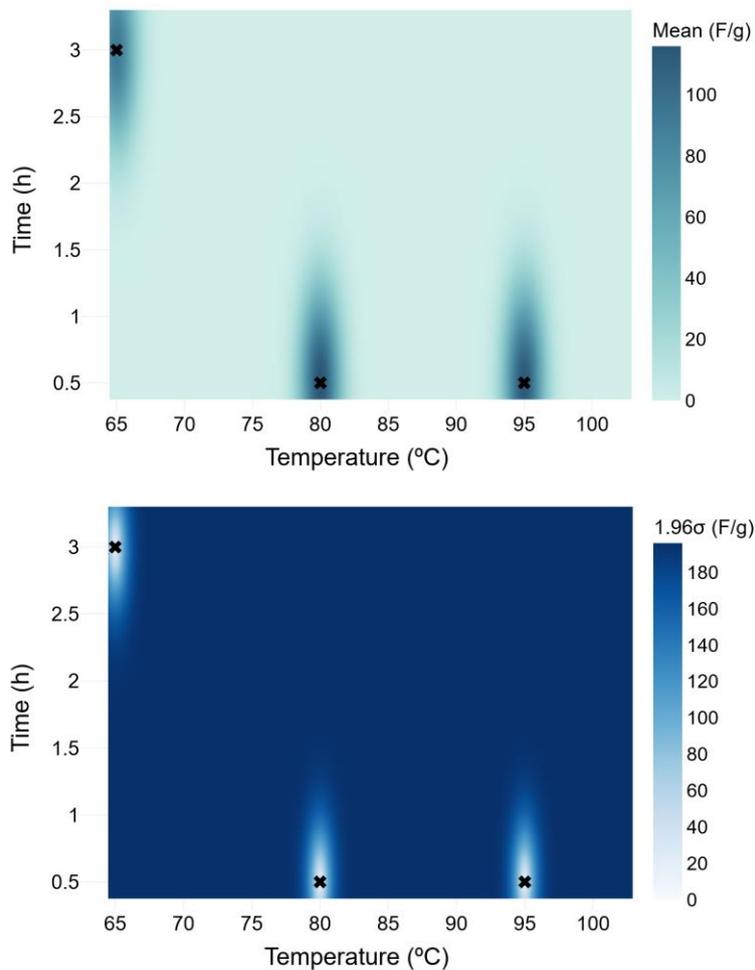


Figure 4-7. Gaussian Processes after three capacitance measurements at 2 A/g, with Temperature (x-axis) and Time (y-axis). The crosses on the plots indicate the points of real data (observations). The top plot shows the predicted mean distribution. The bottom plot illustrates the uncertainty in terms of the 95% confidence interval, calculated as 1.96 times the standard deviation $\sqrt{\sigma_n^2}$, where $\sqrt{\sigma_n^2}$ is computed according to Equation 2-7.

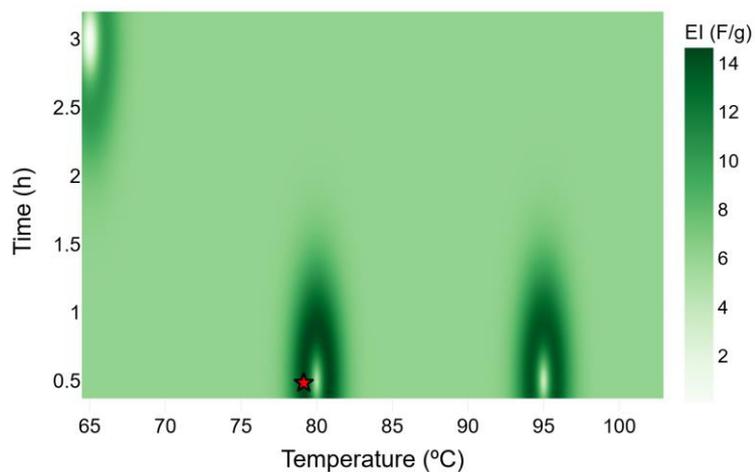


Figure 4-8. Expected Improvement for the Gaussian Process estimated from capacitance values of graphene oxide electrodes reduced at 80°C (0.5h), 65°C (3h) and 95°C (0.5h). The red star indicates the point where the EI reaches its maximum value in the first iteration.

Compared to the initial GP model shown in Figure 4-7, where the posterior mean remained close to zero due to the limited number of observations and the zero-mean prior, the addition of new experimental data during Bayesian Optimisation leads to a rapid upward shift in the GP mean. This transition occurs predominantly during the early iterations of the optimisation, once the model has observed several capacitance values of similar magnitude. Subsequent iterations primarily refine the local structure and uncertainty rather than further shifting the global mean, resulting in a plateau around 100-110 F/g.

The mean capacitance distribution, represented in the top plot of Figure 4-9, remains relatively smooth, fluctuating between 100 and 110 F/g, suggesting a weak dependency of the capacitance on temperature and time within the explored range. The uncertainty heatmap (bottom plot of Figure 4-9) reveals higher uncertainty in regions with fewer training points, particularly in the upper-right and lower-left corners, as well as in the region between 70-80°C above 2 hours. These effects arise partly from the limits of the decision space, as introduced earlier. The GP is defined over a bounded domain of 65-95°C and 0.5-3h. In areas close to these boundaries, especially at high temperatures, the model has fewer neighbouring observations to support its predictions, which causes the GP to revert toward its prior mean and results in larger predictive variances. This indicates that the model is less confident in those areas. While refining the model in these regions could be beneficial, the observed smooth capacitance distribution suggests that further tests within this optimisation framework may yield limited new insights.

In fact, this smooth behaviour motivates the study of other parameters beyond temperature and time, which may have a greater impact on determining capacitance. Examples could include parameters affecting the drying step [409,410], the oxygen content of the initial graphene oxide [411,412] or post-heating treatments [413].

LOO-CV was performed, with residuals presented in Figure 4-10. All held-out residuals fall within their respective 95% confidence intervals. This demonstrates that the model's uncertainty estimates (governed by the noise parameter σ_n and the anisotropic RBF kernel) closely reflect the true variability in the data. The kernel's length-scales were validated as appropriate, with no systematic under/over-prediction across the experimental range.

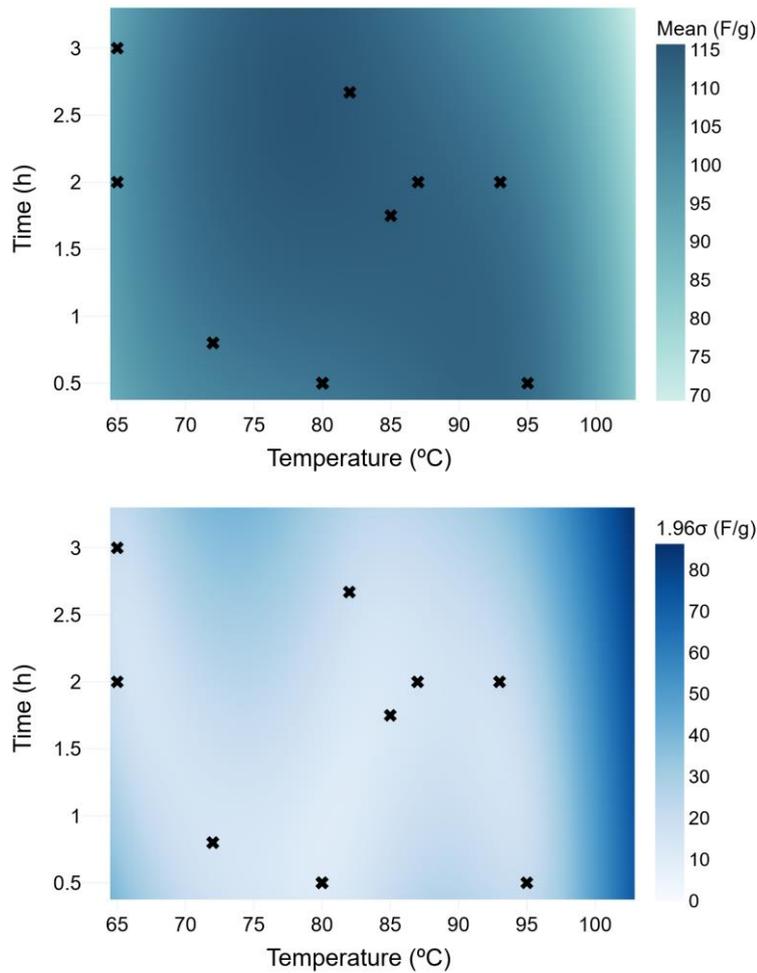


Figure 4-9. GP mean prediction distribution (top) and 95% confidence level (bottom) after 13 observations.

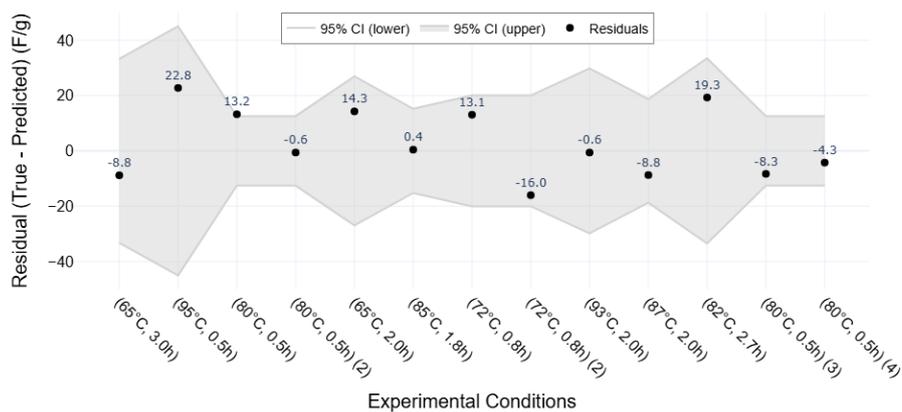


Figure 4-10. Residual plot for each held-out condition during LOO validation.

Previously shown GCD curves and capacitance data indicate a pseudocapacitance behaviour of rGO electrodes, which is expected from this kind of material [411]. While its high surface area and porous structure contribute to Electrical Double Layer Capacitance (EDLC), functional groups such as hydroxyl, carboxyl and carbonyl, that remain present

after an incomplete reduction of the original graphene, can act as active sites for Faradaic redox reactions, causing pseudocapacitance.

Figure 4-11 presents two types of electrochemical characterisations for the reduced graphene oxide electrodes. The top panel shows Nyquist plots, where the real part of the impedance appears on the x-axis and the negative imaginary part on the y-axis. Each electrode is represented by two curves: a solid line measured at 0.8 V and a dash-dotted line measured at 1.2 V. The width of the semicircle reflects the charge-transfer resistance, while the low-frequency tail indicates capacitive behaviour. The bottom panel displays the corresponding cyclic voltammetry curves, in which the current response is plotted against the applied potential for the same voltage limits. These plots allow direct visual comparison of the capacitive and pseudocapacitive contributions under different reduction conditions and operating voltages.

The Nyquist plots reveal that higher voltage causes a reduction of the width of the half-circle at medium-high frequencies, which is attributed to a reduction in the charge transfer resistance R_{CT} [392,414]. This suggests enhanced redox activity at higher voltages, as the applied potential drives faster kinetics. Moreover, the voltage-dependent behaviour at the low-frequency region confirms the pseudocapacitive nature of the electrode, as an ideal EDLC would exhibit vertical lines, independent of the voltage applied. Cyclic voltammetry curves display quasi-rectangular shape with broad, asymmetric features in both the anodic and cathodic branches, confirming the pseudocapacitance behaviour for these materials [415]. In addition, the increase in voltage window is accompanied by a rise in current and capacitance, suggesting the activation of redox processes enabled by the wider voltage window.

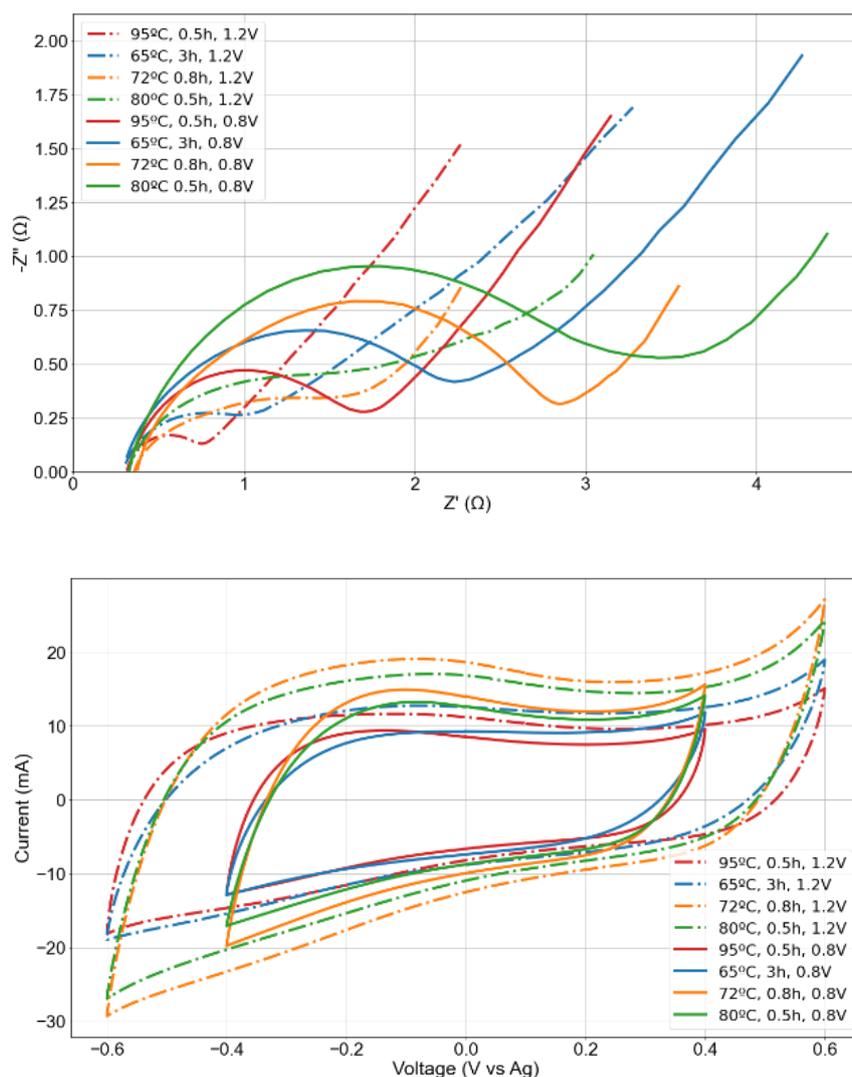


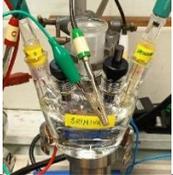
Figure 4-11. Top: Nyquist plots for rGO electrodes reduced under different temperature and time conditions, measured at 0.8V (solid lines) 1.2V (dash-dotted lines). Bottom: corresponding CV curves for the same electrodes and voltage conditions.

4.2.4 Interlaboratory Modelling

To gain deeper insights into the effectiveness of GPs in predicting capacitance from noisy observations, a set of samples was characterized using three-electrode method, with a different set-up and in a different laboratory (i.e., Nth-order replication, as defined in Section 3.1). This set-up is referred to as ‘set-up 2’ to distinguish it from the original configuration, referred to as ‘set-up 1’. The potential sources of uncertainty in a model that covers capacitance values evaluated from both set-ups are varied and include several factors, such as variations in the measurement equipment, different reference and counter electrodes, fluctuations in the electromagnetic interferences, and variations in cell geometry and electrode spacing [416]. Moreover, testing protocols differed slightly; specifically, in set-up 1, the voltage window for the GCD was 1.2V, while in set-up 2, this window was limited to 0.8V. All these differences are intended to develop a model with high generalisation power which integrates the interlaboratory measurement protocol

fluctuations observed in literature. Table 4-5 describes the differences between these two set-ups.

Table 4-5. Characteristics of the two set-ups used to evaluate the capacitance of rGO electrodes.

| | Set-up 1 | Set-up 2 |
|----------------------------|---|---|
| |  |  |
| Cell | 3D-printed Swagelok cell | Gamry's EuroCell™ |
| Current collector | Stainless steel bar | Ni foam |
| Counter Electrode | Activated Carbon Norit SA Super (90%) + PTFE (10%) | Pt sheet |
| Reference Electrode | Ag bar | Ag/AgCl/saturated KCl |
| Electrolyte | 1M KOH | 1M KOH |
| Potentiostat | BioLogic | Gamry |
| Voltage window | 1.2V | 0.8V |

The model was then updated by two means: firstly, by incorporating an additional set-up dependent deviation term to account for systematic differences between the measurement configurations, and secondly, by introducing a new categorical input feature indicating the set-up used. This additional term is not intended to represent random noise but rather the systematic shift or bias introduced when capacitance measurements are performed across different laboratories and instrumentation.

To estimate this deviation associated with the type of set-up, GCD curves of one electrode were measured in set-up 1 and compared with GCD curves from another electrode measured in set-up 2. Both electrodes had similar masses (3.34 mg and 3.4 mg, respectively) and contained rGO from the same reduction process. The GCD curves are represented in Figure 4-12, while Table 4-6 summarises the results.

The GCD curves measured in set-up 1 and set-up 2 clearly show the influence of the available voltage range on the electrochemical response. In set-up 2, the charge-discharge cycles finish earlier because the measurement is constrained to 0.8 V, while in set-up 1 the curves extend further in time due to the larger 1.2 V window. This means that, even at the same applied current density, the electrode in set-up 1 is storing and releasing charge over a wider potential range.

At $1 \text{ A} \cdot \text{g}^{-1}$, the GCD curves from set-up 2 exhibit a more symmetric and triangular profile, with very little curvature. This behaviour is typically associated with ideal double-layer capacitive storage and high coulombic efficiency, since the charging and discharging processes are nearly mirror images of each other. The reduced curvature suggests that

pseudocapacitive or faradaic contributions are less significant when the potential window is limited to 0.8 V.

In contrast, the curves obtained in set-up 1 display slightly more deviation from linearity. This is consistent with the activation of additional redox or pseudocapacitive processes at higher potentials, which increases the apparent capacitance at low current densities (e.g., 117.81 F·g⁻¹ vs 99.13 F·g⁻¹ at 1 A·g⁻¹) but may also reduce the ideal triangular shape.

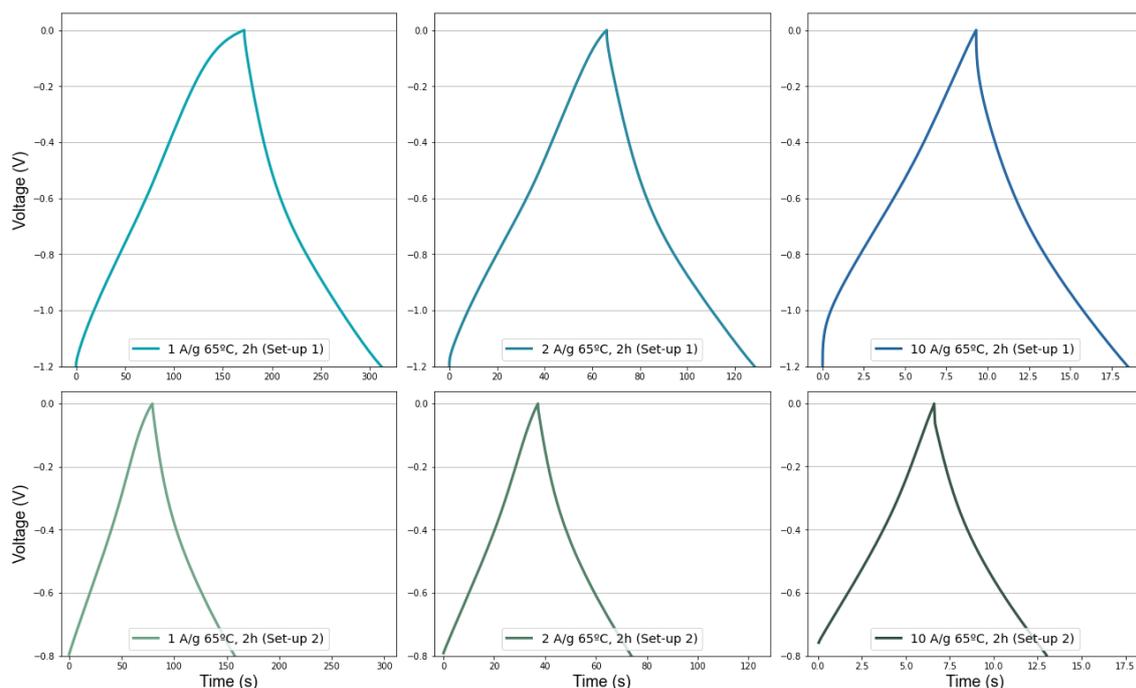


Figure 4-12. GCD curves for two different reduced graphene oxide electrodes measured in two different experimental set-ups in two different labs. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right). Top figures correspond to measurements carried out using set-up 1, bottom figures refer to measures using set-up 2. The graphene in the electrodes was reduced from graphene oxide at 65°C for 2 hours.

Table 4-6. Specific capacitances at 1 A/g, 2 A/g, and 10 A/g for rGO electrodes measured in two different setups.

| Set-up \ Current Density | 1 A/g | 2 A/g | 10 A/g |
|--------------------------|-------------------|-------------------|------------------|
| Set-up 1 | 117.81 F/g | 104.55 F/g | 78.22 F/g |
| Set-up 2 | 99.13 F/g | 93.54 F/g | 86.88 F/g |

Table 4-7. Noise attributed to measurements coming from different setups, for 1A/g, 2A/g and 10A/g.

| | | | |
|----------------------------|------------------|-----------------|-----------------|
| $\sqrt{\sigma_{setups}^2}$ | 13.21 F/g | 7.78 F/g | 6.13 F/g |
|----------------------------|------------------|-----------------|-----------------|

Thus, a total noise has been recalculated, by adding the previous noises related to heat source and mass fluctuations, to the noise originating from differences between set-ups.

$$\sqrt{\sigma_n^2} = \sqrt{\sigma_{heating}^2 + \sigma_{mass}^2 + \sigma_{setups}^2} \quad \text{Equation 4-13}$$

Table 4-8. Total noise considered for 1 A/g, 2 A/g and 10 A/g.

| | 1 A/g | 2 A/g | 10 A/g |
|---------------------|------------------|------------------|------------------|
| $\sqrt{\sigma_n^2}$ | 16.94 F/g | 14.20 F/g | 29.52 F/g |

With this updated total noise, a new variance (σ_n^2) was computed in the White Kernel (Equation 4-7) for the GP that takes the values characterized in set-up 2.

Beyond this noise update, a setup-specific bias was incorporated to account for systematic offsets between measurement systems. The GP kernel was extended to include a categorical feature for setup identity (one-hot encoded as [1,0] or [0,1]) as it follows:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \cdot \exp\left(-\frac{\|\mathbf{x}_i^{[T,t]} - \mathbf{x}_j^{[T,t]}\|^2}{2l_{T,t}^2} - \frac{\|\mathbf{x}_i^{[s]} - \mathbf{x}_j^{[s]}\|^2}{2l_s^2}\right) \quad \text{Equation 4-14}$$

where $\mathbf{x}^{[T,t]}$ are the temperature/time features, $\mathbf{x}^{[s]}$ is the setup feature, and $l_{T,t}^2$ and l_s^2 are the temperature/time and setup length-scales, respectively. Following the same physics-informed philosophy as in the first GP, l_s was fixed by assuming an inter-setup tolerance of $\leq 15\%$. This value was derived from the relative change in capacitance attributable to variations in the voltage window in carbon-based nanostructures, calculated using the decision tree rules of ref. [417], together with the standard deviation associated with data analysis in pseudocapacitors, extracted from the interlaboratory study in ref. [363]. This yields $l_s = 2.3$ (see Appendix B, page 199 for detailed derivation).

Unlike temperature and time, which are controllable process parameters deliberately explored within a constrained but continuous design space defined by scalability and cost considerations, the measurement setup is a categorical and non-controllable variable. Its role in the model is not to describe a tuneable process response, but to account for systematic interlaboratory offsets. For this reason, the $\leq 15\%$ assumption defines a bounded range in the output space (i.e., the maximum expected deviation in capacitance attributable to setup differences), rather than a range of input smoothness to be inferred from data. This tolerance is therefore encoded as a fixed length-scale, preventing the GP from attributing excessive variance to the categorical feature while ensuring physically realistic inter-setup differences.

To assess the performance of this GP, LOO validation was conducted. Because the heatmaps were nearly identical when the interlab data were included, they are omitted here for clarity. Additionally, a data point from [418] was incorporated for testing. This was the only point found in the literature referring to the capacitance of reduced graphene oxide, where the reduction conditions, namely the reducing agent, temperature and time, fell within the limits considered in the present model. For this point, the categorical feature corresponding to set-up 1 was assigned, as the conditions in this setup were the closest to those described in [418]. The 95% confidence interval in Figure 4-13 is computed from the GP predictive distribution as $\hat{y} \pm z_\alpha \sigma_{\text{pred}}$, where \hat{y} is the predictive mean, σ_{pred} is the predictive standard deviation returned by the model, and $z_\alpha \approx 1.96$ for a 95% interval. The noise variance was set to $\alpha = 201.7 \text{ (F/g)}^2$, corresponding to an experimental standard deviation of 14.2 F/g, in line with the uncertainty estimates reported in Table 4-8, so the confidence intervals already reflect the expected measurement errors. Figure 4-13 shows the residuals from both LOO validation and testing.

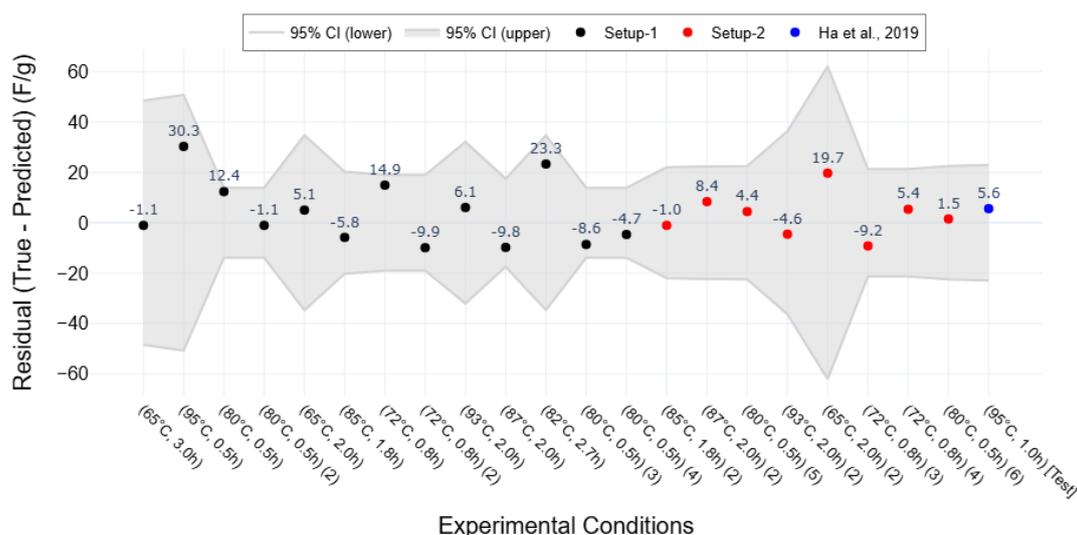


Figure 4-13. Residual plot for each held-out condition during LOO validation. Black dots represent capacitance values measured with Setup-1, while red dots correspond to Setup-2. The blue dot indicates the residual of the test point, whose reference value was taken from ref. [418]. The prediction for this point was computed using the following inputs: 95°C, 1h and Setup-1.

The residuals fluctuate around zero, indicating a mix of underestimations and overestimations without systematic bias. This confirms the appropriate choice of kernel and suggests that the model is sufficiently complex to capture the underlying trends. Furthermore, all residuals lie within the 95% confidence interval, further validating the model's robustness. The testing point, taken from ref. [418], is accurately predicted, with only a minor underestimation of 5.2 F/g. This point serves as an external validation of the model, providing an independent literature value against which its predictive

performance can be assessed. For comparison, Figure 4-14 shows results from a GP model where no categorical feature was added to account for different setup configurations, leading to a setup-induced offset.

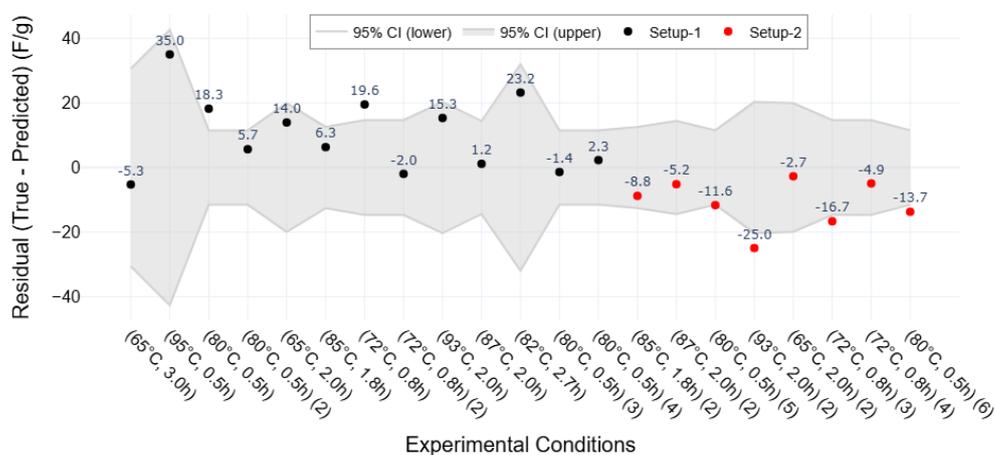


Figure 4-14. Residual plot for each held-out condition during LOO validation under the two setup configurations.

The previously described pseudocapacitive nature of the rGO materials may be the origin of the observed bias between setups in Figure 4-14. The higher voltage window in set-up 1, compared to set-up 2 is potentially activating higher redox groups, causing the observed higher capacitance values.

The three largest deviations (at 95°C, 0.5h; 82°C, 2.7h; and 65°C, 2.0h) occur in regions of higher uncertainty, suggesting that additional data in these ranges could improve the model. Most residuals (19/22) fall within the ± 20 F/g range, which aligns with the most frequent bounds of the 95% confidence intervals. This threshold is derived from the estimated measurement noise (Table 4-8) combined with an irreducible error component. As described before, the measurement noise was obtained by evaluating the variability introduced by selected experimental uncertainty sources. These include temperature fluctuations during the reduction step, differences in electrode mass, and systematic differences between measurement set-ups. For each source, the standard deviation of the capacitance measured from paired replicates was calculated. The three noise contributions were then treated as independent and combined using the quadratic sum, which gives the overall experimental noise level expected across the model domain. The origin of this residual error may be attributed to aleatoric effects such as environmental fluctuations affecting sample preparation or testing, limitations in measurement resolution, or potential biases introduced during data processing. Additionally, epistemic uncertainties could contribute. For example, the unknown spatial distribution of the binder, conductive additive and reduced graphene oxide in the

electrode. A systematic quantification of these uncertainty sources would be valuable for future studies.

4.3 Conclusions and Future Work

In this chapter, a physics-informed Gaussian Process model was developed to guide the synthesis of reduced graphene oxide with improved capacitance, using a non-toxic and scalable reduction method based on ascorbic acid. The result is a heatmap that relates temperature and time conditions to capacitance. To efficiently explore the synthesis space, Bayesian Optimisation was employed, using GP as surrogate model with the motivation of reducing the number of experimental trials required.

Unlike purely data-driven models, the GP developed here incorporates experimentally determined noise terms and applies physical constraints to the kernel length-scales, based on scalability criteria. This grey-box approach ensures that the model respects the underlying physics of the system, accounting for variability arising from reduction conditions, electrode preparation and electrochemical testing, which are all key contributors to the often-cited reproducibility challenges in materials science.

The GP's predictive performance, validated through Leave-One-Out Cross-Validation (LOOCV), confirms appropriate hyperparameter selection, with all experimental points falling within the 95% confidence interval. The resulting mean capacitance values remain relatively stable, ranging from 100 to 110 F/g across the tested range of temperature and time. This smooth behaviour suggests a weak dependency on these variables within the explored window and highlights the need to investigate additional synthesis parameters that may exert stronger influence. Examples include drying procedures, original oxygen content of the graphene oxide or innovative post-treatment processes. Electrochemical characterisation using Nyquist plots and cyclic voltammetry supported the pseudocapacitance behaviour of the rGO electrodes, complementing the model's insights.

From an economic perspective, the reduction of graphene oxide represents one of the most versatile strategies for tuning capacitance, because it allows controlled modification of porosity and surface functionalities such as oxygen groups. The method also benefits from starting with graphene oxide, a material already produced at industrial scale in quantities of several thousand tonnes per year, which significantly improves cost-effectiveness and scalability. In this work, the selection of reduction parameters, including temperature, reaction time and the use of an inexpensive and non-toxic reductant, was guided by the goal of minimising the economic impact of the process while maintaining electrochemical performance. These considerations support the

potential of this approach to achieve a competitive performance-to-cost ratio when compared with other electrode materials used in supercapacitors, such as activated carbons, carbon nanotubes and doped carbons.

To evaluate generalisability across experimental setups, a second interlaboratory data set was introduced, which is important because models developed from single-laboratory data often fail when applied to different experimental environments. Variations in equipment, electrode preparation, measurement protocols and operator practices can introduce systematic differences that significantly affect electrochemical properties. Assessing performance across independent data sets therefore provides a more rigorous test of the model's robustness and its potential for broader applicability.

The GP was updated with an extended noise term and an additional kernel component to capture systematic biases between setups. A 15% inter-setup tolerance was adopted, motivated by prior studies highlighting capacitance variability in pseudocapacitors. The updated model continued to perform well, as confirmed by LOOCV and the successful prediction of an external test point from the literature. Together, these results underscore the value of combining ML with physical insight to enhance understanding and predictive control in the synthesis of advanced electrode materials, while contributing to improved reproducibility in experimental materials science and supporting further research critical to the transition toward a zero-emissions economy.

There are several potential directions for continuing the work presented in this chapter. One of these involves refining the feature engineering process, which in this study yielded input features with limited influence on capacitance, likely due to the constraints imposed for scalability. Exploring heat treatments after the chemical reduction could be a promising extension of the current model.

Regarding uncertainty analysis, it would be valuable to investigate the influence of current density on the experimentally observed noise in greater detail. This may require a more comprehensive evaluation of experimental variability (e.g. including multiple samples per condition), as well as deeper electrochemical analysis to uncover underlying patterns that are affected by current density. Additionally, exploring the stability of electrodes during prolonged cycling, and how this is influenced by synthesis input parameters, offers another meaningful research direction.

More broadly, the uncertainty estimation framework described in this work could be extended to support research on other materials, such as emerging battery cathode candidates, or even in fields beyond electrochemistry such as additive manufacturing of new alloys, membrane development, or bioplastics.

5 Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity

Summary

The rise of Machine Learning models has impacted all scientific domains, from biology and chemistry to physics and materials science. In materials science, while these models offer clear advantages in research, they are sometimes constructed without ensuring physical correctness. Moreover, conventional performance metrics often fail to adequately assess the physicochemical validity of the complex patterns identified by these models. Additionally, the trade-off between material structure specificity, key chemical and physical properties, and data availability presents a challenge in training models for specific material families. Achieving sufficient complexity to capture underlying patterns while avoiding overfitting remains a key difficulty.

In this work, a ML model for graphene capacitance is constructed using a limited dataset. A series of steps are introduced during model creation that, for this specific scientific domain, are innovative. Such steps include outlier identification using Mahalanobis distances of feature distributions across publications, grouped k-fold cross validation where the groups are composed of datapoints from the same publication, and introducing a metric that penalizes error variability between those groups to find the best hyperparameters. All these measures have the motivation of guaranteeing physical correctness while optimizing the model complexity and avoiding overfitting. Such correctness is deeply assessed with literature examples from outside the dataset. According to the constructed model, the top three main physicochemical features affecting graphene's capacitance are nitrogen doping, oxygen doping and average pore diameter. These principles for ML model development can be extended to other applications in materials research. A final comparative evaluation with experimentally synthesised samples confirmed the complementary strengths of GP and RF models, highlighting their joint value for rGO electrode design.

5.1 Introduction

As introduced in section 2.3, a ML model is a mathematical descriptor of a given collection of structured data (dataset) that adjust the weights or parameters to minimise the errors of such description. The effectiveness of an ML model in either case is influenced by the nature of the model, defined by its inputs, outputs, the quantity and quality of points in the dataset and the algorithm followed to fit the model to the data [342,419].

The use of ML in the field of electrochemistry has become a prolific area of research. Key applications include the discovery of new electrocatalysts for energy storage, harvesting and conversion, as well as the production of sustainable chemicals [420–433]. ML also supports the interpretation of electrochemical data by identifying patterns and relationships in measurements such as voltammetry, impedance spectra and cycling behaviour that are difficult to extract using conventional analysis [434–444], and enables advanced diagnostic and prognostic techniques in electrochemical devices [445–455]. As the field continues to evolve, researchers are constantly exploring and developing additional innovative applications of ML in electrochemistry.

Before the rise of ML models in electrochemical research, traditional approaches to model creation consisted of the ideation of a series of equations, supported by the established knowledge from experimental evidence, previous models or a combination of both on the relevant fields at the time the model is created. This physical-chemical laws aim to represent in an unambiguous manner how a part of the material main characteristics explains some of its properties under a given characterisation context (e.g. Gouy-Chapman model, Butler-Volmer equations, etc.).

In contrast, the approach of an ML scientist aiming to create a new model differs from traditional equation-driven methods, as it relies on data-driven learning rather than physical-chemical laws [342]. Knowledge of the physical or chemical fields governing the system under study is not necessarily required to be explicitly encoded in order to obtain a model that achieves reasonably accurate predictions according to predefined accuracy and/or error metrics. Depending on the algorithm selected, this model may provide somehow ‘understandable’ relationships or, on the contrary, a set of given conditions or equations whose interpretation is unaffordable for a human mind.

As a consequence of this, some current approaches may lack proper consideration of material science and electrochemical fundamentals, both in their construction and in the interpretation of results. This can lead to models that, while statistically robust, may not adhere to the principle of ‘physical correctness’, i.e., the alignment with established physicochemical laws and principles. Consequently, some models may contradict each other regarding which physicochemical characteristics are most relevant or may inappropriately equate materials of different natures or measurements made under disparate conditions [456–462]. Examples of these contradictions in the context of graphene capacitance are described in section 1.3.3.

In this chapter, this problem is tackled in the context of ML models for predicting capacitance in carbon-based material, with a focus on graphene and its derivatives. The

aim is to develop a novel ML that adheres to the principles of physical correctness while effectively mitigating overfitting.

5.2 Methods

5.2.1 Feature Selection

The objective of the model presented here is to gain a deeper understanding of the relative importance of graphene's physicochemical features on capacitance, according to the current literature. The model's inputs consist of physicochemical characteristics that have been established in the literature as key factors influencing capacitance. Table 5-1 below summarises them, with an explanation of their technical relevance as indicators of capacitance, and references from other carbon materials models that included them as inputs.

Table 5-1. Technical relevance and presence in previous model of the selected physicochemical features of graphene.

| Input variable | Technical relevance | Presence in previous capacitance models | | | | | | | | | | | |
|--------------------------------------|---|---|-------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|-------|-------|-------|
| | | [282] | [283] | [284] | [285] | [286] | [287] | [290] | [295] | [298] | [303] | [304] | [305] |
| Oxygen % | Contribute to pseudocapacitance | | X | | X | X | | X | X | X | X | X | X |
| Sulphur % | Contribute to pseudocapacitance | | | | | X | | | | X | X | X | |
| Nitrogen % | Contribute to pseudocapacitance and electrical conductivity | X | X | | X ^b | X | | X | X | X ^b | X | X | X |
| Specific Surface Area | Contribute to EDL capacitance | X | X | X ^a | X ^a | X ^c | X ^d | X ^a | X ^d | X | X | X | X |
| Total Pore Volume | | | | X | | | X ^c | X ^d | | X ^d | | | |
| Average Pore Diameter | Condition the EDL capacitance and ion transport | X | X | | | X | X | | X | | | | |
| I _D /I _G ratio | Condition the EDL capacitance and electron transport | X | X | | | X | | | X | X | X | X | X |

^aIn this model, the surface area contribution is divided into two indicators: one for the surface area derived from micropores (SA_{micro}) and another for mesoporous (SA_{meso})

^bNitrogen % is divided into 4 indicators: pyrrolic nitrogen, pyridinic nitrogen, quaternary nitrogen and other nitrogen type

^cBoth surface area and pore volume are each split into two indicators: one for the total surface area/pore volume and another for the relative proportion corresponding to micropores

^dBoth surface area and pore volume are each divided into three indicators: one for the micropores (S_{micro} , V_{micro}), one for the mesopores (S_{meso} , V_{meso}) and one for the total pore contribution (SSA, PV)

The model is intended for samples where capacitance has been characterised using Galvanostatic Charge-Discharge measurements under KOH electrolyte. This technique and electrolyte, as opposed to Cyclic Voltammetry and other electrolytes, represent the most commonly used configurations in previous models. Although the influence of the electrolyte can be important and interesting to assess, the current model is focused on the most widely adopted setup to avoid having excessive inputs, thereby mitigating the risk of the curse of dimensionality [463].

Likewise, no distinction was made regarding electrolyte concentration, testing method or voltage window, as data collection was limited to publications using three-electrode setup with 6M KOH and a 1 V voltage window, which is the most frequently reported configuration for these materials. Furthermore, the durability during cycling has not been addressed in the model, restricting the output to the capacitance obtained after the first cycles.

Following a similar philosophy of maximising the spectrum of available data in the literature while maintaining a moderate number of inputs, no distinction has been made between contributions from micropores and mesopores for either surface area or pore volume. Similarly, nitrogen content has been considered as a single input. These choices constitute deliberate simplifying assumptions introduced to limit input dimensionality and to ensure consistency with datasets and input definitions employed in previously published models of graphene-based electrode capacitance reviewed in this work. Although pore-size-resolved descriptors and detailed nitrogen functionalities can provide additional insight into charge storage mechanisms, incorporating them would require substantial data curation and redefinition of inputs across studies, significantly reducing the number of usable samples. As a result, the present model is intended to capture dominant, average trends across the dataset rather than to explicitly capture effects associated with pore size or nitrogen chemistry.

Finally, current density is also included as input variable. The reasons behind this are two-fold: first, by including the current density as input, the number of available data from the same kind of material is maximised, as very frequently the same sample is characterized at different current densities. Secondly, current density and capacitance are typically negatively correlated, regardless of the physicochemical features of the sample. So, by including it as an input, this underlying contribution is being considered, apart from the contributions from physicochemical properties.

With respect to data balance, the dataset is not statistically balanced in the sense of providing uniform coverage of the entire input space. The degree and nature of imbalance vary across variables. Capacitance and oxygen content exhibit relatively broad and unimodal distributions, providing good coverage of the ranges most commonly reported in the literature. Other inputs, including current density, heteroatom content, specific surface area, pore volume, average pore width, and the I_D/I_G ratio, display varying degrees of skewness or long-tailed behaviour. Importantly, for several of these variables the skewness is centred around physically meaningful and frequently reported values (for instance, I_D/I_G ratios close to unity in graphene-based carbons), indicating

that the observed imbalance reflects realistic experimental practice rather than a sampling artifact. A detailed discussion of the distributions of pore volume and average pore width is provided later in this thesis, while the remaining input distributions are reported in Appendix C. This imbalance is inherent to literature-derived datasets and was intentionally preserved to ensure representativeness rather than artificial uniformity. As the present work addresses a regression problem rather than classification, strict balance across input variables is not a prerequisite. Consequently, model predictions are expected to be most reliable within densely populated regions of the input space, while extrapolation toward sparsely sampled extremes should be interpreted with caution.

5.2.2 Data Collection

Once the inputs of the model are defined, the next step is to collect data containing information about those inputs. As previously mentioned, the data is restricted to characterisations performed using the GCD technique in a KOH electrolyte for graphene-related materials. It is important to note that, in the context of supercapacitors, graphene-related materials refer to the following variations, as defined in ref. [464]: few-layer graphene (FLG), multi-layer graphene (MLG), graphene oxide (GO), and reduced graphene oxide (rGO).

Thus, active carbon, amorphous carbon, carbon nanotubes, carbon black and graphene composites have been excluded. This exclusion is intended to focus on the goal of understanding what are the relative importances of physicochemical features for the family of graphene materials. As discussed in the section 1.3.3, current models in the literature do not seem to address this clearly. For this reason, data gathering focused on datasets from previous models that include involving graphene materials. After filtering the initial dataset, which consisted of 673 papers and 1540 datapoints, and based on the considerations introduced in the previous section related to the electrochemical technique, material structure, and electrolyte, the final dataset contained 257 datapoints derived from 23 papers, as individual studies typically reported multiple measurements under varying experimental conditions.

5.2.3 Data Preprocessing

The dataset has been pre-processed with the goals of improving the quality of the datapoints and adjusting the outliers to facilitate model training. This was carried out by first completing the missing values and secondly by identifying and filtering outliers.

Different strategies can be employed to handle missing values in datasets. One of these involves replacing missing values with zeros. However, this approach can negatively affect the physical correctness of the model, especially when applied to variables such

as SSA, pore volume and diameter or I_D/I_G . Then, this strategy has been followed only for the specific case of data related to chemical composition that was missing (mostly sulphur content). In ref. [298] authors completed the missing I_D/I_G values by imputing the arithmetic mean of the available values. However, as previously stated, this approach can have negative consequences, especially when comparing the importance of features relative to each other, as is the intention here. First, the artificial imputation of the mean for certain instances reduces the variance for that input, thereby affecting the underlying information provided by other points in the literature are providing and diminishing its statistical power. Second, it also impacts the relationships between this input and others, distorting those relationships. Thus, in our case, the presence of I_D/I_G was a *sine qua non* condition for retaining the instance in the dataset; consequently, all entries lacking a reported or directly extractable I_D/I_G value were excluded during the data curation process.

Most of the missing values in our dataset corresponded to average pore diameters. For the cases where the pore distribution curves were available in the original publications, an estimation of this feature was carried out by first extracting the points, then using Akima interpolation [465] to approximate the curve and finally calculating the area behind the curve using Simpson's rule to extract the average pore diameter (Figure C1 shows a diagram describing the procedure followed). This method was also used to replace the data where the estimations had been made using the following relationship

$$Pore\ size = 4000 \cdot \frac{Pore\ volume}{Specific\ Surface\ Area} \quad \text{Equation 5-1}$$

This relationship is an approximation from Wheeler's theory that assumes a homogeneous and uniform distribution of cylindrical pores. Such an assumption can lead to a loss of variance associated with this input variable [466].

The next step involved filtering the data to remove outliers, with a primary focus on outliers in the input space. Such outliers correspond to regions of low data density or to isolated studies that are difficult to generalize from, as their influence may reflect study-specific effects rather than underlying process behaviour. First, a method for defining outliers is required (e.g., IQR in [295] or Hat indices in ref. [286]). Then, a policy is applied to exclude these regions (e.g. excluding data outside 1.5 times the IQR), ensuring a more robust and interpretable model.

To identify outlier publications based on their input characteristics, this work used the Mahalanobis distance to compare the distributions of input features across studies. Each

publication was represented by an aggregated and normalized vector of its input variables, effectively summarizing the relative representation of each feature within that study.

The code is available here [367]. By computing pairwise Mahalanobis distances between these vectors, it was possible to assess how similar or dissimilar each paper was relative to the rest of the dataset, accounting for correlations between features. The average distance of each paper to all others was then used as a proxy for its representativity within the sampled input space: studies with consistently high average distances were identified as statistical outliers in terms of representativity, that is, as statistically isolated with respect to the bulk of the data. Importantly, this classification does not imply that such studies are physically incorrect or unrealistic, but rather that they operate in sparsely sampled regions of the input space. As a result, their inclusion makes it difficult to distinguish whether observed effects were due to the underlying process or from study-specific design choices, potentially reducing the model's generalisability.

To the best of my knowledge, this criterion has not been previously applied in the context of capacitance models or similar ones. However, it is hypothesized that this approach shows promise for ML models in nanomaterial science, where data is derived from experimentation. In the field of nanomaterials research, experimental reproducibility is often complicated by intrinsic variability and statistical noise arising from multiple factors, including differences in experimental protocols, reagents workflows, human factor and equipment.

This issue can be even more pronounced in the present context, where a relatively small dataset is compiled from data originating from multiple laboratories, each introducing its own unique 'noise signature' (as discussed in Section 4.2.4). With this approach, the aim is to identify a subset of input variables that are relatively homogeneous across data sources, minimising the influence of laboratory-specific noise and creating a more controlled and consistent space for model training. Figure 5-1 shows the Mahalanobis distances calculated between the 23 publications used to build the dataset, together with the distribution of these distances (see Table C1 in Appendix C for references). The left panel shows the distribution of pointwise distances across all publications, while the right panel presents the deviation of each publication's mean Mahalanobis distance from the global mean. This figure is important because it quantifies the consistency of each publication relative to the overall feature distribution. Publications with larger mean distances tend to report materials or testing conditions that deviate from the majority of the literature, which may also correlate with different measurement conventions or

synthesis practices. The analysis therefore provides a data-driven way to identify potential inter-study heterogeneity, which later helps explain variability in predictive-model performance.

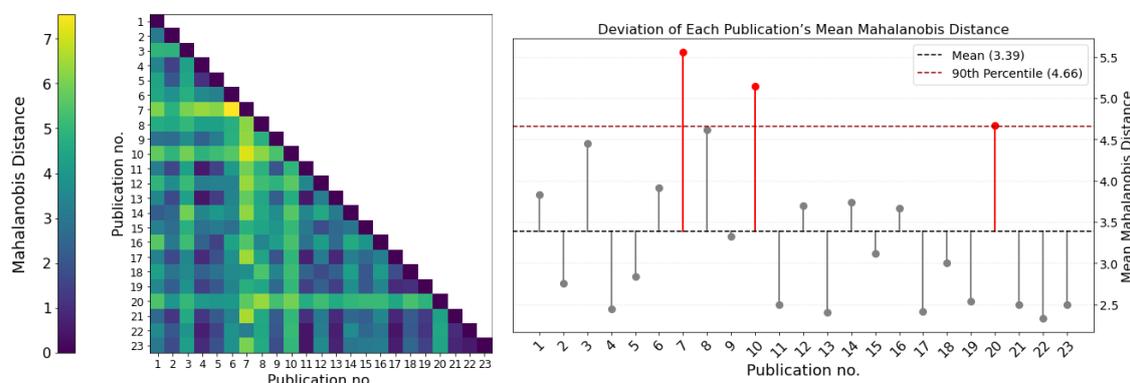


Figure 5-1. Left: Mahalanobis distances across publications from which datapoints were extracted to build the dataset. Right: Deviation of each publication's mean Mahalanobis distance.

Publications with distributions ranked in the top 10% of Mahalanobis distances, defined as those exceeding the 90th percentile (indicated by the dark red dashed line in the right plot of Figure 5-1), were classified as outliers and subsequently removed. This threshold was selected as a conservative criterion to identify only the most extreme multivariate observations while retaining the majority of the dataset, thereby minimising the risk of discarding representative samples. Accordingly, these outliers correspond to publications¹ numbered 7, 10 and 20 in the left plot of Figure 5-1. Their classification as outliers may be explained by distinct features of the data: for instance, publication no. 20 reports unusually high pore widths (~ 80 nm compared to an average of ~ 8 nm across the other publications), publication no. 10 shows the largest pore volume ($164 \text{ cm}^3 \cdot \text{g}^{-1}$ versus an average of $\sim 0.72 \text{ cm}^3 \cdot \text{g}^{-1}$), and publication no. 7 combines one of the lowest specific surface areas with one of the top three highest pore widths. Following this, a second filtering step was implemented to identify outliers in the capacitance data. Specifically, any data points falling outside 1.5 times the interquartile range (IQR) were classified as outliers. As a result, 1 point was removed from each of the publications 6, 8 and 18; 2 points from both publications 12 and 21; and 4 points from publication² 16. It should be emphasised that these data points were removed based on a statistical criterion aimed at improving robustness and representativity, rather than due to any assessment of their experimental validity

¹ References [499], [502] and [512] in the bibliography, respectively

² References [498], [500], [510], [504], [513] and [508] in the bibliography, respectively

5.2.4 Model Selection

Six models were evaluated using the filtered dataset: Linear Regression, Support Vector Machine, Decision Tree Regressor, Artificial Neural Network (ANN), Random Forest, and Gradient Boosting. Optimal hyperparameters were determined using Grid Search K-fold Cross-Validation with $k = 5$. The train/validation splits were organized by grouping the data based on the publications from which each subset was extracted. The reasons for this grouping are explained below. Except for the ANN model, which was built using TensorFlow [467] and follows a similar architecture to that in ref. [298], but adapted to eight inputs, all other models were constructed using the Scikit-learn package [468].

Many of the publications considered in the initial dataset include data from the same sample evaluated at different current densities. They may also contain samples with varying physicochemical properties synthesized by the same research group, under a similar laboratory environment and technical circumstances. A frequent mistake in ML models for capacitance prediction is treating these subsets as entirely independent samples during model training. Consequently, models are often trained and evaluated on instances exhibiting similar 'group-induced variance' and/or the same sample characterized at different current densities is present in both the train/test splits. This practice can cause model overfitting and undermine statistical robustness and generalizability. Thus, grouping the data by publication during cross-validation tasks aims to mitigate these issues and ensure statistical validity.

In addition, a validation metric is proposed in this work, to incorporate variability across groups during hyperparameter selection (see Equation 5-2 below).

$$vpRMSE = RMSE + \lambda \sigma_{RMSE} \quad \text{Equation 5-2}$$

$vpRMSE$ stands for variance-penalized RMSE, σ_{RMSE} is the standard deviation of RMSE across different train/validation splits, and λ is a weighting factor that controls the trade-off between overall accuracy and robustness in the optimisation. This metric was used to select the best hyperparameters for each model using Grid Search. The idea behind this is to ensure robustness and generalisation power for the model. This term accounts for the variability of results across different publications during validation and ensures robustness and generalisation capability. The results presented below correspond to an optimisation with $\lambda = 0.5$.

Although it is common in ML-based capacitance models to associate good model performance with an R^2 value close to one, this approach can be counterproductive, especially when considering models with a certain degree of specificity and a limited

dataset size, as it is the case in this study. This is because R^2 represents the ratio of variability explained by the model to the total variability in the data. In other words, if the dataset contains noise, as is the case here with experimental data collected from different laboratories, a model that fits this noise more closely may yield a higher R^2 value. This increases model complexity while limiting its ability to identify general patterns within the constrained dataset, ultimately promoting overfitting [469].

While information criteria such as the Akaike Information Criterion (AIC) [470] and the Bayesian Information Criterion (BIC) [471] explicitly penalise model complexity and are commonly used for parametric model comparison, they were not adopted in this work for several reasons. First, the primary objective of the present study is not formal model selection among competing probabilistic models, but the identification of relationships between physicochemical features and capacitance that generalise across studies, in line with previous ML studies in this field. In this context, the literature on capacitance prediction using ML models predominantly relies on error metrics combined with cross-validation, rather than likelihood-based information criteria, and methodological consistency was considered important in order to allow comparison with earlier work.

Second, several of the models evaluated in this study, including Random Forests, Gradient Boosting and Artificial Neural Networks, are non-parametric or semi-parametric. For these models, the definition of an explicit likelihood function, and consequently of AIC or BIC, is not straightforward and would require additional assumptions that are difficult to justify when working with experimental data collected across multiple publications.

Finally, given the limited size of the dataset and the variability between publications, model evaluation was deliberately based on cross-validation schemes that separate data by publication and on performance metrics that assess predictive behaviour on unseen studies. Within this framework, error metrics derived from cross-validation are more closely aligned with the objectives of this study than information criteria. Nevertheless, AIC and BIC may represent useful complementary tools in future work, particularly as larger and more consistent datasets become available or when the analysis is restricted to fully parametric models. The performances of the models were assessed using Leave-One-Publication-Out Cross-Validation (LOPOCV), adhering to the principle of excluding data from the same publication in both training and testing.

In summary, model selection was performed using a nested cross-validation strategy that respected publication-based grouping. Specifically, the outer loop used Leave-One-Publication-Out Cross-Validation (LOPOCV) across 19 unique publications.

For each fold, data from one publication was held out entirely for testing, while the remaining 18 publications formed the training set. Within this training set, the inner loop applied grouped 5-fold cross-validation (i.e., folds were split by publication) to optimise hyperparameters. This ensures that, at every stage, no data from the held-out publication was used in either training or validation, preserving a strict separation between model selection and evaluation. This structure avoids data leakage and provides a reliable estimate of model generalisation across unseen publications.

5.3 Results

Figure 5-2 shows the real vs predicted capacitance values for each of the testing folds during LOPOCV for the five models. The observed dispersion reflects the combined effects of inter-publication variability and the Leave-One-Publication-Out Cross-Validation strategy, rather than a lack of predictive structure. In this context, the figure is intended to illustrate consistency of predictive behaviour across heterogeneous studies, while quantitative model comparison is performed using aggregated error metrics computed consistently across folds.

Table 5-2 summarises the performance metrics of them as its mean \pm standard deviations across folds. Both the figure and the table are intended to assess generalisation across heterogeneous studies rather than pointwise prediction accuracy. The best RMSE values were achieved by Random Forest, followed by Linear Regression. The strong performance of Random Forest is consistent with its ensemble nature, which effectively reduces overfitting by averaging predictions across decorrelated trees, improving both robustness and accuracy. Linear Regression performs better than more complex models likely because the dataset relationships are partially linear, allowing it to capture key trends despite its simplicity. The observed preference for the linear kernel over the RBF kernel in SVM model during Grid Search further suggests the partially linear nature of the dataset. However, the overlapping intervals with other models make definitive conclusions difficult.

Across all models, both RMSE and vp-RMSE values show broadly comparable performance between test and validation sets, with most differences falling within ranges of statistical variability. While some models (e.g., Linear Regression, ANN) exhibit slightly higher validation error than test error in RMSE, this relationship reverses uniformly for

vp-RMSE, where all models show lower test values than validation. This systematic pattern likely arises from how the penalty term interacts with the different data aggregation strategies (single-paper in test versus multi-paper in validation sets). However, just as with standard RMSE, the overlapping standard deviations in vp-RMSE suggest that these systematic differences are not statistically significant.

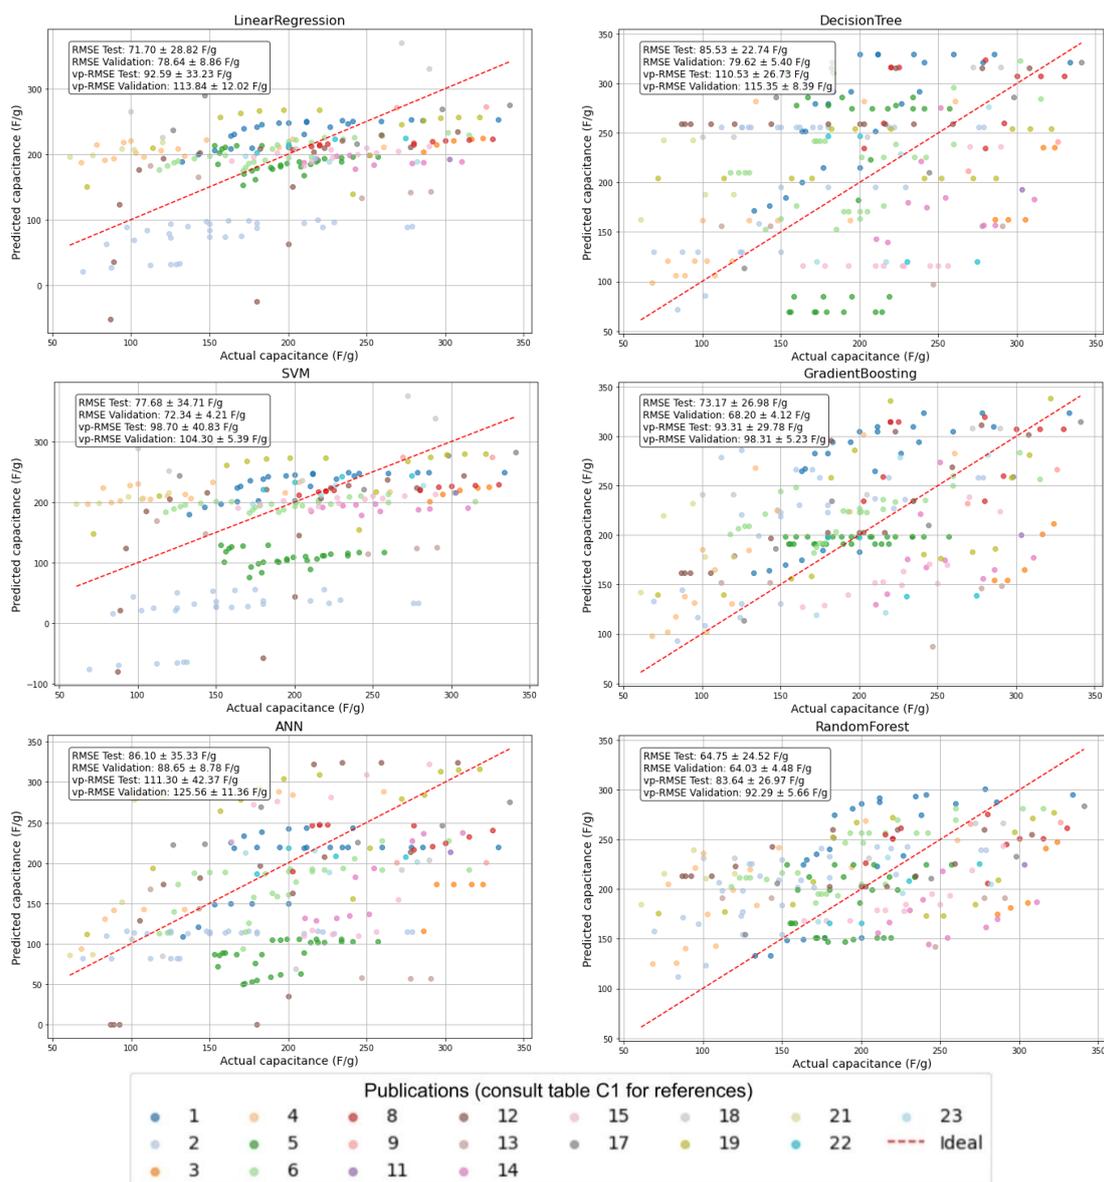


Figure 5-2. Real versus predicted capacitance values for each testing fold under Leave-One-Publication-Out Cross-Validation. The dispersion reflects both model error and inter-publication variability inherent to literature-derived experimental data. The figure illustrates consistency of predictive behaviour across heterogeneous studies rather than pointwise prediction accuracy.

Table 5-2. Summary of performance metrics for the five models during validation and testing.

| Model | RMSE Test (F/g) | RMSE Validation (F/g) | vp-RMSE Test (F/g) | vp-RMSE Validation (F/g) |
|------------------|-----------------|-----------------------|--------------------|--------------------------|
| LinearRegression | 71.70 ± 28.82 | 78.64 ± 8.86 | 92.59 ± 33.23 | 113.84 ± 12.02 |
| SVM | 77.68 ± 34.71 | 72.34 ± 4.21 | 98.70 ± 40.83 | 104.30 ± 5.39 |
| DecisionTree | 85.53 ± 22.74 | 79.62 ± 5.40 | 110.53 ± 26.73 | 115.35 ± 8.39 |
| GradientBoosting | 73.17 ± 26.98 | 68.20 ± 4.12 | 93.31 ± 29.78 | 98.31 ± 5.23 |
| ANN | 86.10 ± 35.33 | 88.65 ± 8.78 | 111.30 ± 42.37 | 125.56 ± 11.36 |
| RandomForest | 64.75 ± 24.52 | 64.03 ± 4.48 | 83.64 ± 26.97 | 92.29 ± 5.66 |

At first sight, the magnitude of the prediction errors may appear large. However, these values must be interpreted in the context of a literature-derived dataset that combines measurements from multiple laboratories, synthesis routes, and testing protocols. In this setting, prediction error reflects not only model uncertainty but also experimental variability that is not encoded in the input features. The consistency between validation and test errors across models, together with the comparable performance of linear and non-linear methods, indicates the presence of systematic relationships between the selected physicochemical descriptors and capacitance, which are masked by significant inter-study noise. This point is examined in detail in the discussion that follows.

Percentage-based error metrics were not adopted, as they tend to amplify errors at low capacitance values and are not commonly reported in comparable ML models for capacitance prediction. Instead, absolute error metrics were preferred, as they allow direct comparison across models and better reflect the practical uncertainty associated with predicting capacitance from heterogeneous experimental data.

The higher standard deviations in test metrics (e.g. ± 28.82 F/g for Linear Regression RMSE test, ± 8.86 F/g for RMSE validation) can be attributed to the LOPO strategy used in the outer loop of nested cross-validation. Since each test fold represents data from a single paper, inter-paper variability likely amplifies the observed fluctuations. In contrast, the inner-loop validation metrics (calculated via grouped k-fold CV with multiple papers per fold) benefit from aggregated data, reducing variability and producing more stable estimates.

Random Forest achieves the lowest RMSE values in both test (64.75 ± 24.52 F/g) and validation (64.03 ± 4.48 F/g). Overfitting was assessed by examining the consistency between validation and Leave-One-Publication-Out test errors under publication grouped cross validation, together with the variability of these errors across folds. The close agreement between validation and test RMSE, combined with a comparatively low

variance-penalized RMSE, indicates limited overfitting relative to the other non-linear models. Although the wide test interval reflects sensitivity to inter-paper variability inherent to literature-derived datasets, it does not correspond to a systematic degradation in predictive performance on unseen publications. In contrast, Artificial Neural Networks and Decision Trees exhibit larger dispersion in test RMSE and higher variance-penalised RMSE values, such as the 28% increase observed for ANN, indicating stronger sensitivity to data splits and publication-specific biases.

Although the obtained values for Random Forest are higher compared to those from other models in the literature (Table 1-7), the metrics here have been calculated to truly reflect the model's ability to generalise capacitance results, which are assumed to have publication-induced bias.

To illustrate the effect of ignoring publication-based data grouping, the performance of the same Random Forest model, with identical hyperparameters ($\text{max_depth}=4$ and $\text{n_estimators}=16$), was also assessed using a conventional random 80-20% train-test of the full dataset. The upper subfigure in Figure 5-3 shows the real versus predicted capacitance values obtained from one such random split, whereas the lower subfigure corresponds to a second, independently generated random split using the same procedure.

At first sight, the performance metrics obtained for the first split ($\text{RMSE}=34.20$ F/g and $\text{R}^2=0.77$) appear superior. However, these results are misleading because, when the data are split randomly without grouping by publication, samples from the same studies appear in both the training and test sets. This allows the model to partially 'recognise' patterns it has already seen, artificially inflating the apparent predictive performance.

Moreover, the comparison between the two subfigures highlights that the resulting metrics change noticeably when a different random split is used. Although the visual distribution of the predicted values remains broadly similar in both cases, the variations in RMSE and R^2 demonstrate that this approach is highly sensitive to the specific way the data are partitioned. This limits its reliability when working with literature-derived datasets. In contrast, the publication-grouped split provides a more rigorous and realistic evaluation framework, because a model that performs well under this scheme is more likely to generalise to entirely new studies.

Figure 5-4 shows the SHAP (SHapley Additive exPlanations) summary plot, based on the Random Forest Regressor trained with the whole dataset. SHAP is an additive feature attribution method grounded in cooperative game theory, where each input feature is treated as a 'player' contributing to the model prediction. For a given instance,

SHAP computes feature contributions by averaging their marginal impact on the prediction over all possible feature combinations, resulting in a locally accurate and consistent decomposition of the model output.

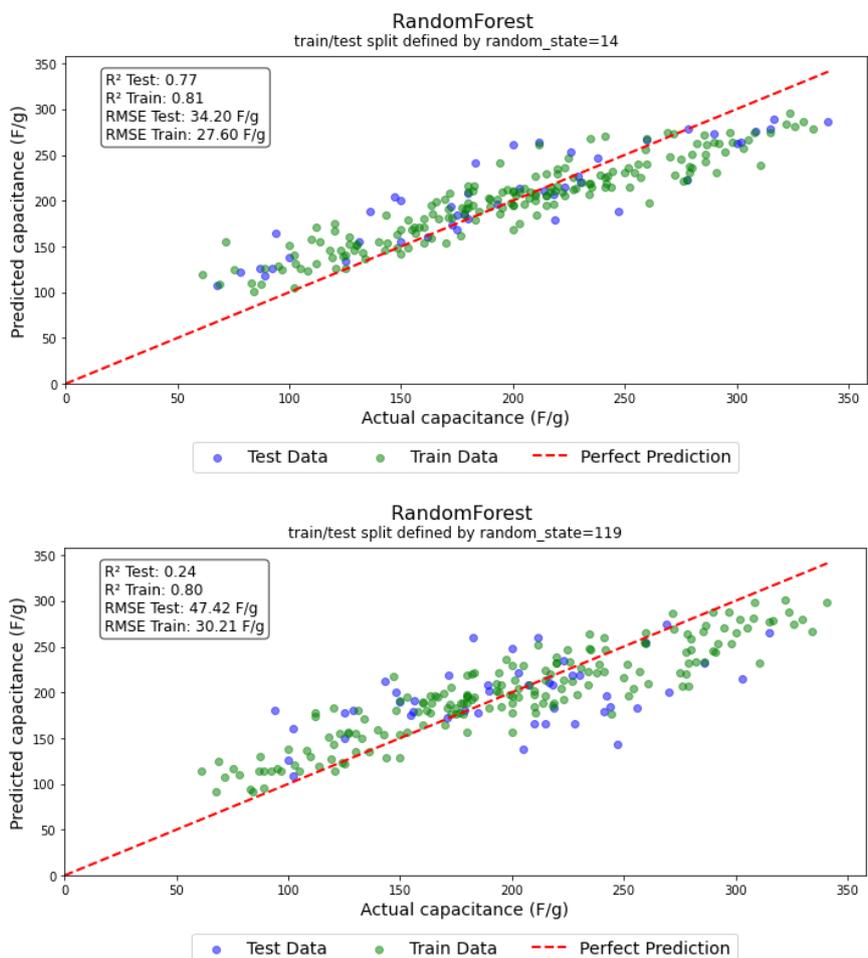


Figure 5-3. Real versus predicted capacitance values for the Random Forest model evaluated using two different random 80-20% train-test splits of the full dataset. Although the visual distributions appear similar, the resulting performance metrics differ, reflecting sensitivity to data partitioning.

Unlike traditional global sensitivity analysis, which quantifies average input-output relationships across the entire dataset, SHAP provides instance-level explanations that can be aggregated to obtain global feature importance while accounting for feature interactions [472]. In this study, SHAP was used to interpret the output of the ML model by assigning importance values to each feature, and the TreeExplainer algorithm was employed due to its computational efficiency and suitability for tree-based models such as Random Forests [473].

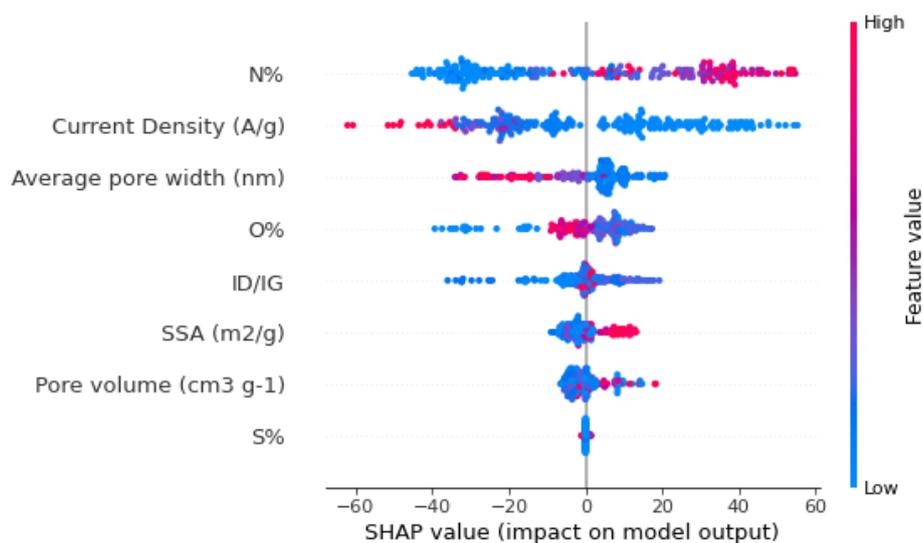


Figure 5-4. Summary plot of SHAP values demonstrating the influence of features on the Random Forest Regressor model.

In the summary plot, the features are ranked vertically by their overall importance, with the most important feature at the top. This importance is calculated as the mean absolute value of the SHAP values for each feature across all instances. A SHAP value represents the marginal contribution of a given feature to the model prediction for a specific instance, computed by evaluating its contribution under all possible feature combinations. Each point on the plot represents a SHAP value for a specific feature and a specific instance. The x-axis shows the SHAP value, which indicates how much the feature contributed to pushing the model's output higher or lower for that instance. As extracted from the lateral colour bar, the colour of each point represents the value of the feature for that instance. Thus, among all the physicochemical characteristics of reduced graphene oxide, nitrogen content is the most important according to this analysis, with high values of it as the main driver to increase graphene's capacitance, followed by average pore width, oxygen level, I_D/I_G , SSA, pore volume, and sulphur content.

Since feature importance is an intrinsic property of the model rather than an external performance metric, it is evaluated using the entire dataset. To ensure stability despite the randomness inherent of Random Forests, which is controlled by the `random_state` hyperparameter, the results report the average SHAP values over 100 independently trained models, each initialised using a different `random_state`, to stabilize feature importance across different model initialisations. Figure 5-5 illustrates the mean absolute differences between consecutive cumulative means of the SHAP values for all features as the number of training iterations increases. The rapid decrease in these differences during the initial iterations indicates that feature importance estimates are

highly sensitive when based on only a few model realisations. As the number of iterations increases, the differences progressively approach zero, demonstrating convergence of the SHAP values and stability of the feature importance rankings. The near-zero differences observed beyond approximately 40-50 iterations confirm that averaging over 100 trainings provides robust and reproducible SHAP-based feature importance.

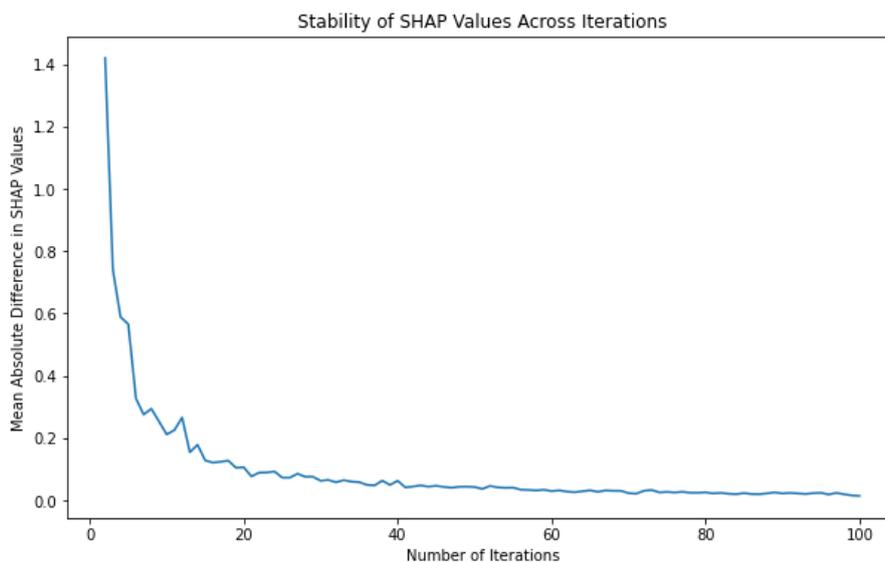


Figure 5-5. Mean Absolute Differences between consecutive cumulative means of the SHAP values for all the features after 100 iterations.

5.3.1 Physical Correctness Assessment

The dominant influence of nitrogen over other physiochemical features such as average pore size, oxygen or SSA, as identified by the SHAP analysis, aligns with findings from publications not included in the dataset, supporting the reliability of the model. Śliwak *et al.*[474] compared the capacitance of non-doped reduced graphene oxide electrodes with 16.2 O at.% and SSA of $530 \text{ m}^2\cdot\text{g}^{-1}$, with a nitrogen-doped (N-doped) counterpart with 9.4 O at.%, 13.4 N at.% and an SSA of $354 \text{ m}^2\cdot\text{g}^{-1}$, and using a three electrode configuration and 6 M KOH electrolyte. The N-doped sample, despite having lower SSA and oxygen content, exhibited almost 50% higher capacitance at $20 \text{ A}\cdot\text{g}^{-1}$ and experienced 21% less decay in capacitance after 5000 cycles. This effect was also observed in other electrolytes. For instance, Li *et al.*[475] showed that in 1M TEABF₄ N-doped samples with $677 \text{ m}^2\cdot\text{g}^{-1}$ reached more than 30% higher capacitance than non-doped counterparts which have a surface area of $763 \text{ m}^2\cdot\text{g}^{-1}$. Moreover, the N-doped samples showed a lower average pore diameter, which is consistent with the SHAP analysis indicating that lower pore diameters contribute positively to the predicted capacitance.

The SHAP analysis ranks current density as the second most influential variable, with predominantly negative SHAP values indicating a systematic decrease in predicted capacitance as current density increases. This model-derived trend is well supported by experimental evidence reported in the literature, where reduced capacitance at high current densities is commonly attributed to ion diffusion limitations within the porous structure of graphene-based electrodes [476–480].

The average pore width is, after the current density, the following feature in order of importance in our model. Figure 5-6 shows the distributions of average pore width (top) and total pore volume (bottom) across the dataset, together with key descriptive statistics. The ranges of pore widths in the data span from approximately 0.5 nm to 30 nm. However, the mean, median and mode are 7.57 nm, 4.01 nm and 3.42 nm, respectively, indicating that most samples are concentrated toward the lower end of the range. This indicates that the dataset is predominantly composed of graphene samples with small average pore sizes, highlighting a prevalence of microporous and small-sized mesoporous structures over larger mesoporous and macroporous structures. Additional feature distributions are provided in Figure C2. In comparison to the summary plot, the SHAP dependence plot zooms in on a single feature to show how its values influence the model output. Figure 5-7 displays the SHAP values as a function of the average pore width (x-axis), while the colour scale represents a second interacting feature. It is important to note that this figure is fully derived from the SHAP analysis of the trained model and does not represent an independent correlation or regression, but rather the internal logic of the model.

The trend observed in the dependence plot suggests the importance of an optimal hierarchically distributed porous structure for maximising capacitance. At very low pore widths (<2 nm), the SHAP contributions are generally positive, confirming the beneficial role of micropores, which provide the highest capacitance contribution per unit mass. However, the largest positive SHAP values are not associated with the smallest widths. Instead, the largest positive SHAP values appear around ~2.5 nm, a region that corresponds to the highest pore volume (as indicated by the colour scale). Above this maximum, in the range between ~2.5 and 5 nm, the SHAP values remain relatively high, even though the pore volume is lower, which may indicate that mesopores in this range improve electrolyte accessibility to micropores. For average pore widths larger than ~5 nm, the SHAP values progressively become negative, indicating that larger mesopores and macropores contribute less favourably, possibly due to the reduced proportion of micropores, which offer the highest charge storage per unit mass. Within the apparent

optimal range ($\sim 2.5\text{-}5\text{ nm}$), a wide spread of SHAP contributions is observed, from near zero up to almost $20\text{ F}\cdot\text{g}^{-1}$.

These trends, identified by the model directly from experimental data, are consistent with previous ML studies (refs. [284] and [285]), which report that excessively high micropore surface area can have a detrimental effect on capacitance. They also align with independent experimental findings showing that the highest capacitance is achieved when there is a compromise between maximising the presence of small pores and maintaining efficient ion transport [476,478,481–483].

Although the SHAP dependence plot suggests a physically meaningful trend, the interpretation of the average pore width feature must still be made with caution. The apparent influence of pore width on capacitance is likely to be context-dependent, because this feature compresses the entire pore-size distribution into a single scalar value. As a result, materials with completely different distribution shapes may share a similar average pore width while exhibiting very different electrochemical behaviour. This limits the extent to which clear design rules can be derived solely from this feature.

For similar reasons, it is also difficult to extract a simple monotonic trend from the SHAP dependence plot for pore volume. The pore-volume distribution is strongly right-skewed (Figure 5-6), with the majority of samples lying between approximately 0.1 and $0.8\text{ cm}^3\cdot\text{g}^{-1}$. Within this statistically well-represented range, the SHAP values are generally small and predominantly negative (Figure 5-8), indicating that increasing pore volume typically leads to a slight decrease in the predicted capacitance. This suggests that, for most materials in the dataset, creating additional pore volume does not translate into improved electrochemical performance. A plausible explanation is that a higher pore volume fraction may be accompanied by a reduction in framework density and electrical connectivity, or by the development of larger pores that provide less surface area per unit mass. As a result, the double-layer capacitance normalised by mass decreases compared with structures dominated by smaller pores.

Large positive SHAP values are mainly observed at the extremes of the pore-volume distribution (very low or very high pore volume), but these regions are sparsely populated and should therefore be interpreted with caution.

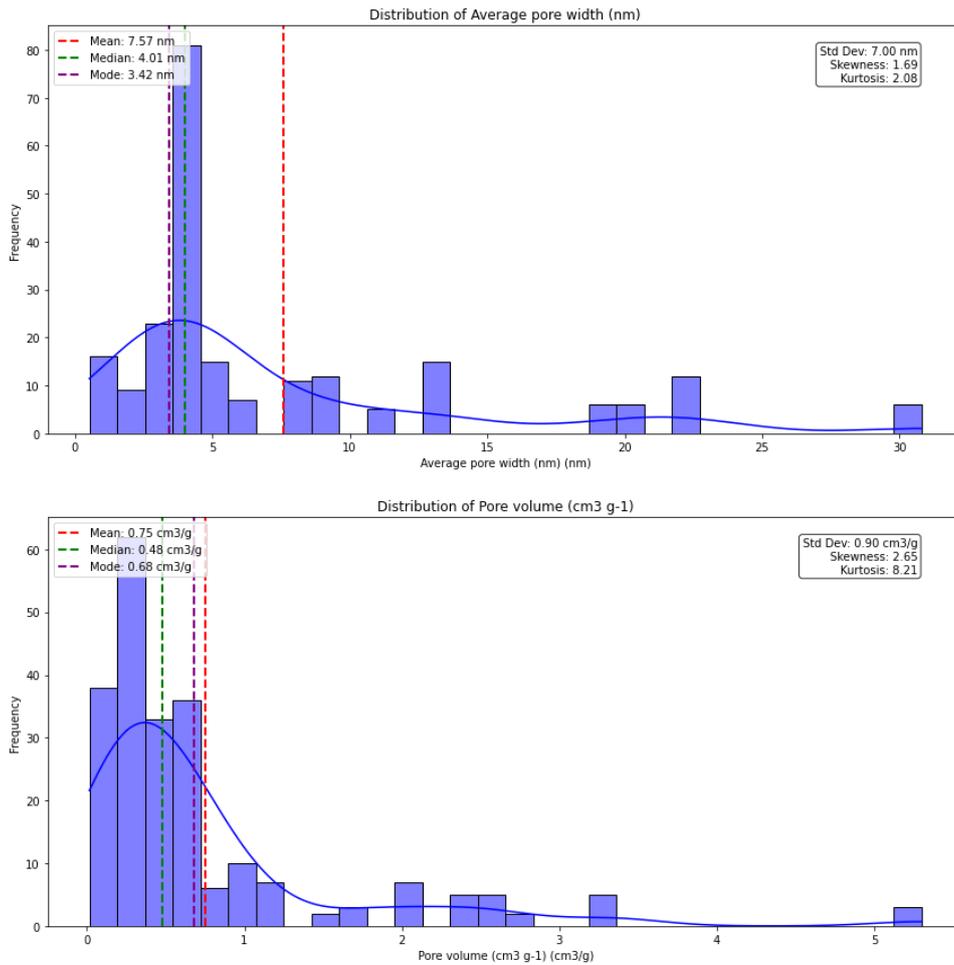


Figure 5-6. Distributions of Average pore width (top) and Pore volume (bottom) with key metrics.

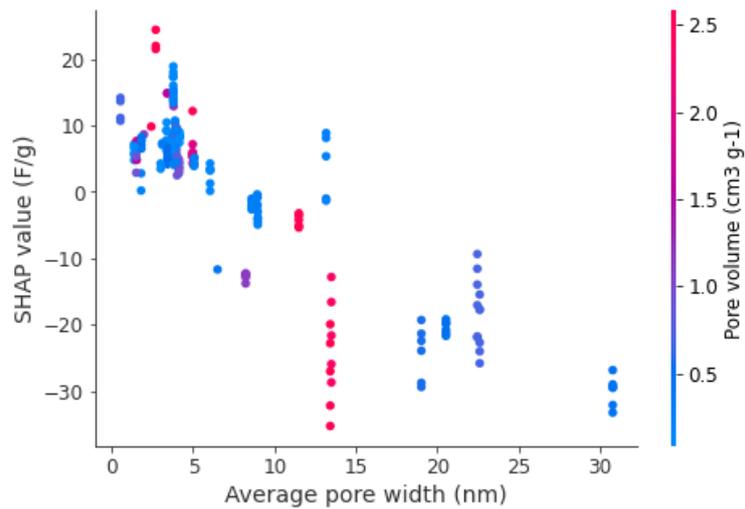


Figure 5-7. SHAP Dependence plot for average pore width, with colours referring to pore volume values.

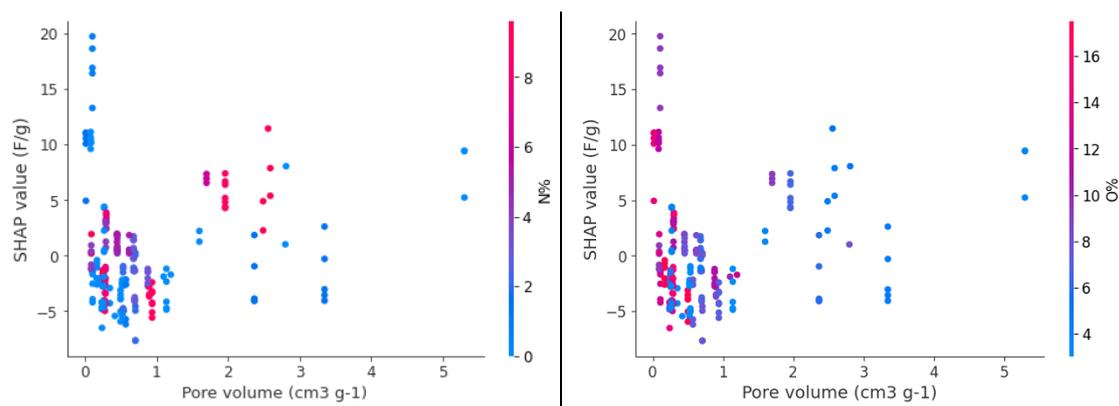


Figure 5-8. SHAP Dependence Plot for Pore volume vs N% (left) and O% (right).

Notably, the strongly positive SHAP contributions tend to coincide with samples containing elevated nitrogen or oxygen contents, as indicated by the colour-coded dependence plots. This suggests that pore volume becomes beneficial primarily when coupled with a high density of surface functional groups, which can enhance wettability and introduce additional faradaic or pseudocapacitive contributions. Therefore, pore volume alone is not a dominant predictor of performance; its effect is strongly modulated by surface chemistry and functional-group concentration.

In respect to oxygen, while in some cases low oxygen levels strongly reduce the capacitance, in others, low oxygen content is associated with positive SHAP values (Figure 5-4). The accumulation of red points at the centre of the x-axis of Figure 5-4 for this feature reveals neutral contribution to the capacitance (neither positive nor negative) when it achieves high levels (~20 O at.%). Intermediate oxygen levels (purple points) are concentrated at positive SHAP values. This behaviour suggests that the role of oxygen in determining capacitance is strongly context-dependent and influenced by other features, such as nitrogen content or structural properties.

This context dependence is further highlighted in Figure 5-9, which shows the SHAP dependence plot of oxygen content with colour indicating nitrogen. Oxygen levels that coincide with high nitrogen contents exhibit clearly positive SHAP values and are therefore beneficial to capacitance.

The work from Yu *et al.* [484], which was not included in the dataset used to train our model, emphasises the importance of the context in respect to the contribution of oxygen to the capacitance in reduced graphene oxide. Samples named NG2 with 8.56 N at.%, 12.88 O at.% achieved approximately 45% and 115% higher capacitance at 1 A·g⁻¹ and 20 A·g⁻¹, respectively, than NG3 samples with 2.11 N at.%, 17.45 O at.%.

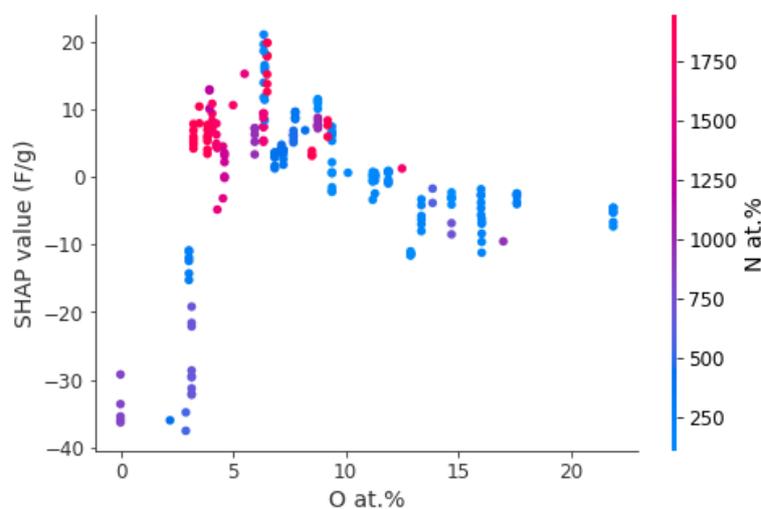


Figure 5-9. SHAP Dependence Plot for O% vs N%.

On the other hand, NG3 outperformed other samples prepared following a distinct functionalisation route, which yielded less oxygen content. Concretely, they showed 40% higher capacitance at $1 \text{ A} \cdot \text{g}^{-1}$ compared to the less-oxidised samples (NGU), which have 11.8 N at.% and 8.21 O at.%. While quantitative information about the structure of the materials was not provided, the authors highlighted the importance of its loose structure to explain the superior performance NG2 samples. This is an example that less oxygen can be comparatively better or worse, depending on the values of other features such as nitrogen or structure features like pore size or specific surface area.

The work of Morimoto *et al.* [485], also not included in the dataset used to train the present model, provides an experimental illustration of this effect but in this case related to the specific surface area. In this work, graphene oxide samples with different oxygen levels were synthesised following two routes: controlled-oxidation of graphite (oGO) and controlled reduction of highly oxidised graphene oxides (rGO). The latter samples showed superior surface area in a wide oxygen range, from ~ 15 O at.% to ~ 50 O at.%. They also showed much higher capacitance for oxygen levels between 10 at.% to 35 at.%, with a peak at 25 at.%.

Figure 5-10 shows the SHAP dependence plot for SSA, with oxygen content encoded in colour. Despite the fact that the data from ref. [485] were not included in our dataset, the linear relationship observed in Figure 5-10 is consistent with the trends reported by Morimoto *et al.* over a wide oxygen range. The existence of a peak at 25 O at.% supports the conclusion that intermediate oxygen levels are optimal to achieve higher capacitance values, compatible with the insights of our model. The poor performance at higher oxygen levels may be due too high charge transfer resistance coming from the low conductivity of the samples, evidenced in the electrical conductivity vs oxygen level plot.

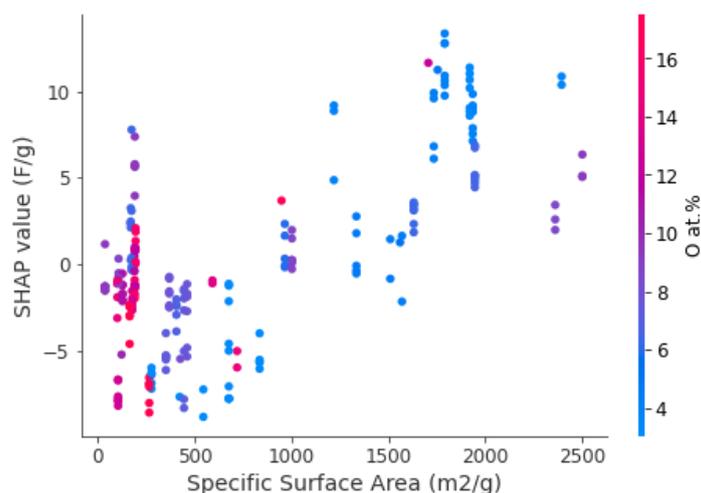


Figure 5-10. SHAP dependence plot of SSA, with colour indicating oxygen values.

Gadipelli *et al.* [411] studied the influence of the initial oxygen content in graphene oxide precursors on the capacitance of rGO samples that were reduced in a preheated vertical tube furnace at 300°C for 2-5 min (EG-named samples). The study also shows the influence of a graphitisation treatment consisting of an annealing for 6h in a horizontal tube furnace under nitrogen gas flow at different temperatures between 400 and 1100°C (GEG-samples). After the first reduction step, the samples coming from graphene oxides with different oxygen levels from 34 to 27 at.% resulted in samples with oxygen levels close to ~12 at.%.

While the differences in oxygen levels across samples were not very large (ranging from 11.13 ± 0.2 at.% to 13.66 ± 0.4 at.%), samples derived from precursors with the highest initial oxygen content exhibited a 200% increase in surface area compared to those from lower-oxygen precursors. Thus, in these samples, the higher capacitance may be attributed to the combination of moderately higher oxygen content and a much larger surface area. Previous studies [481–483,486] have reported the influence of initial oxygen content on the final SSA of reduced graphene oxide samples subjected to rapid thermal shocking.

In ref. [411] the sample named EG5 exhibited ~10% lower oxygen content than EG4 but ~10% higher capacitance, likely due to its lower pseudocapacitance being compensated by a higher EDLC, as it has almost 20% greater surface area. After further annealing steps, which transformed the mildly reduced samples into moderately (~8 at.%) and highly reduced ones (~2.5 at.%), a threefold decrease in capacitance was observed, demonstrating a dominant effect of the oxygen on capacitance. As stated by the authors, aside from the contribution of oxygen functionalities to pseudocapacitance, they can also enhance hydrophilicity/wettability and electrolyte ion adsorption, improving charge storage. However, they also noted that excessively high oxygen concentrations could be

detrimental due to increased charge transfer resistance, resulting from the poor conductivity of sp^3 -type graphene carbon components in oxidised species.

According to our model, the SHAP dependence plot for Specific Surface Area (SSA) reveals two distinct positive linear trends: an initial sharp increase in SHAP values for SSA below $\sim 200 \text{ m}^2\cdot\text{g}^{-1}$, strongly associated with high oxygen content, and a more gradual positive trend at higher SSA values. This suggests that oxygen-rich materials can achieve high capacitance even with modest surface areas likely due to pseudocapacitive contributions from redox-active oxygen functionalities, compensating limited contributions from surface-driven EDLC (Figure 5-10). A similar trend was also reported in ref. [411] where mildly reduced ($\sim 8 \text{ O at.}\%$) and highly reduced ($\sim 2.5 \text{ O at.}\%$) samples exhibited a positive monotonic relationship between SSA and capacitance. This observation further supports the context-dependent role of oxygen and is consistent with the behaviour captured by our SHAP model.

Figure 5-11 presents a Spearman correlation matrix, which provides a global, model-agnostic assessment of monotonic pairwise relationships between the input features and capacitance. The Spearman coefficient (ρ) measures the strength and direction of rank-based associations, but it does not account for nonlinear effects or interactions between variables. The results show that nitrogen content exhibits the strongest positive correlation with capacitance ($\rho = 0.42$), in agreement with the SHAP analysis, which identified nitrogen as the most influential feature with a predominantly positive contribution. Similarly, current density shows a moderate negative correlation with capacitance ($\rho = -0.33$), consistent with the negative SHAP contributions observed across the dataset.

In contrast, several structural and textural features, such as specific surface area, pore volume, and average pore width, display weak direct Spearman correlations with capacitance, despite being identified as important contributors by the SHAP analysis.

This apparent discrepancy arises because SHAP captures multivariate and context-dependent effects, including feature interactions and nonlinearities, which are not accessible through pairwise correlation analysis. For instance, oxygen content shows a strong negative correlation with both SSA and pore volume, indicating that oxygen indirectly influences capacitance through its effect on textural properties. This suggests that more extensive reduction events, which are associated with lower final oxygen levels, may promote an increase in pore presence and SSA during the synthesis of reduced graphene oxide.

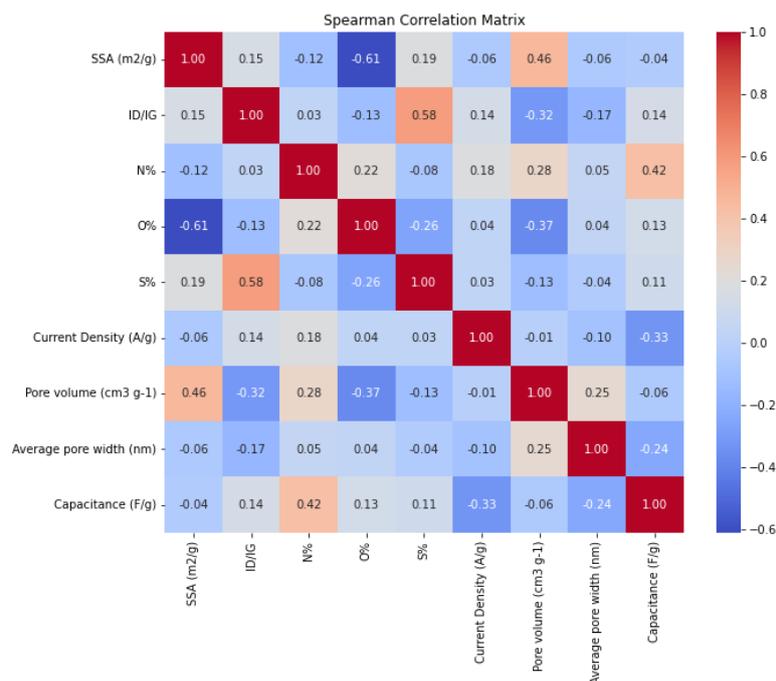


Figure 5-11. Spearman correlation matrix showing relationships between model features and capacitance. While our dataset does not allow for a direct comparison of precursor compositions across studies and considering the specific influence of the thermal shocking treatment in previous findings [481–483,486], our results suggest that lower final oxygen content, potentially indicative of more extensive oxygen reduction events, also correlates with increased SSA, indicating a possible connection between these features in graphene materials.

These interaction-driven relationships are reflected in the SHAP dependence plots but are not fully revealed by Spearman correlation alone. Overall, the consistency in key trends, particularly the dominant role of nitrogen and the negative influence of current density, combined with the complementary insights provided by SHAP and Spearman analyses, reinforces the robustness of the structure-property relationships identified in this work.

Figure 5-12 shows the SHAP dependence of the I_D/I_G ratio on capacitance, together with its interaction with nitrogen, oxygen, and sulphur content. In general, I_D/I_G ratios below 1 are associated with negative or near-zero SHAP values, indicating little or detrimental contribution to capacitance. Most values between 1.1 and 1.3 have positive SHAP values, and it is within this region where this feature has the maximum contribution to the capacitance. A higher I_D/I_G means higher presence of defects, which can be related to holes in the basal plane and/or doping elements.

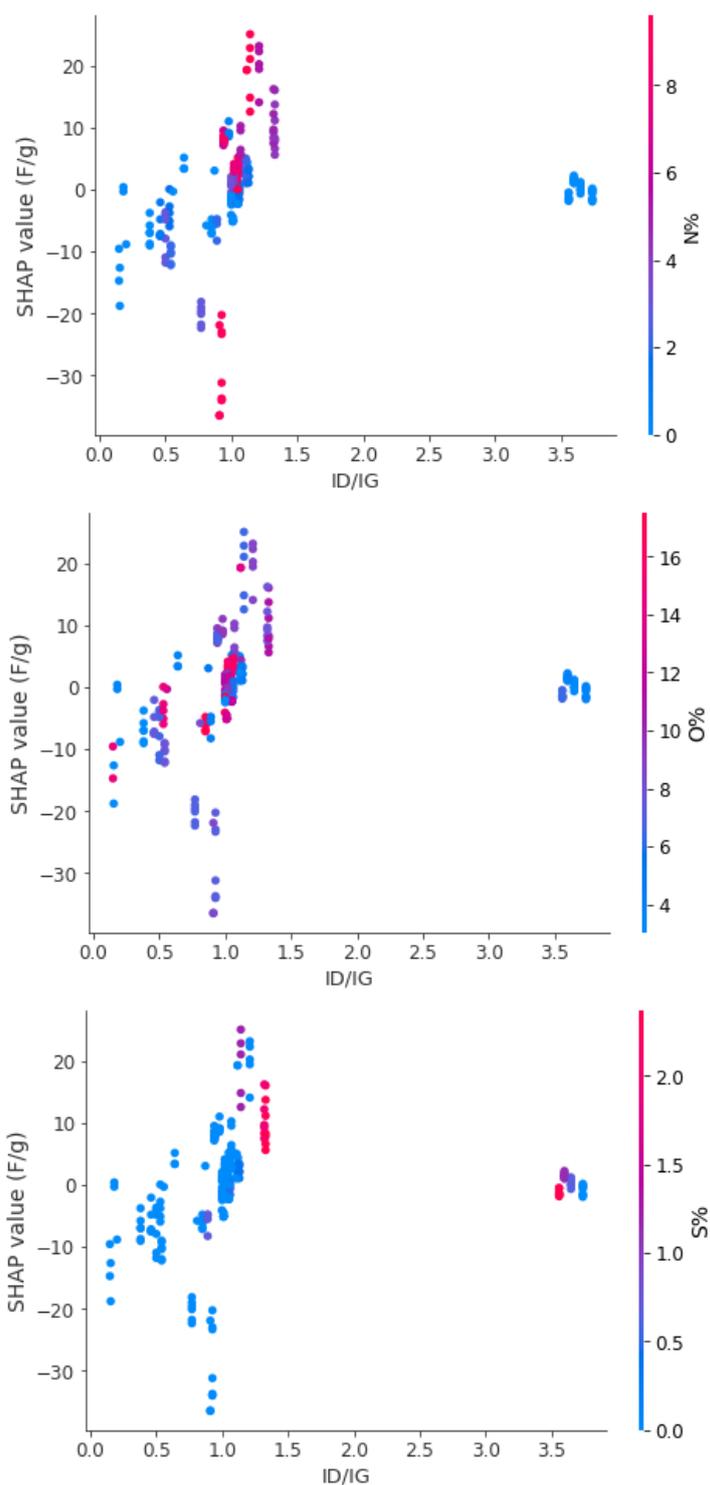


Figure 5-12. SHAP Dependence Plots for I_D/I_G ratio vs. N, O, and S content (top, middle and bottom, respectively).

Both facts can contribute positively to the capacitance, which may explain the more abundant presence of positive values in this region. In fact, as extracted from the three subplots of Figure 5-12, this region of I_D/I_G corresponds to reduced graphene oxide with high levels either of nitrogen, sulphur and/or oxygen, suggesting the influence of

the chemical composition on the effect of the Raman peak ratio to the capacitance. In fact, top SHAP values come when one of the elements has high value (relative to its distribution, indicated by red dots) while the other two have medium values (purple dots).

Lowest I_D/I_G SHAP values match with low sulphur content, as expected from the positive and strong Spearman correlation between these two features (Figure 5-11). The outlier points in the I_D/I_G distribution located at ratios around 3.5, with SHAP values close to zero, suggest that the positive effect of doping may be hindered by the too defective structure, that can compromise electron conductivity [316].

Finally, the summary plot of SHAP values shows that sulphur is the least influential feature; however, the model presented in this work reveals a general positive contribution which reaches its maximum at around 1.2 at.%, as illustrated in Figure 5-13, which represents the SHAP dependence of capacitance on sulphur content. The positive impact of sulphur groups on capacitance has been studied in several works [487–490] and stated in previous ML models for capacitance [286,303,304]. These studies attribute the positive contribution to the role of sulphur groups in enhancing pseudocapacitance and electrical conductivity. While the model developed here indicates a positive contribution from sulphur, the limited representation of sulphur-doped graphene in the dataset (25%) may lead to an underestimation of its relative importance compared to other physicochemical features.

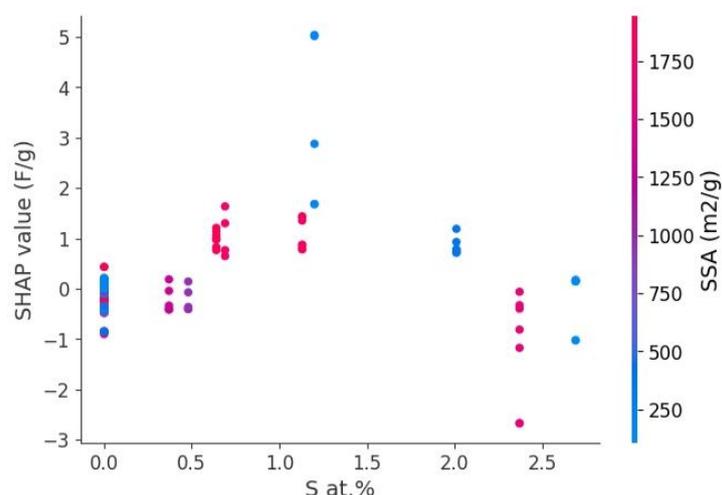


Figure 5-13. SHAP dependence plot of S%, with colour indicating Specific Surface Area values.

5.3.2 Comparative Evaluation of GP and RF Predictions

A final comparative evaluation between the two models developed in this thesis is presented. To this end, the physicochemical properties of a subset of the samples synthesised in Chapter 4 were characterised. Table 5-3 summarises the

physicochemical characterisation of the selected samples. Both nitrogen and sulphur contents were consistently negligible across all samples and were accordingly represented as zero in the RF model inputs.

It is important to note that the purpose of this comparison is not to evaluate model performance in terms of statistical testing metrics (e.g. R^2), which have already been established separately for each model. The GP model was validated using leave-one-out cross-validation on experimental data in Chapter 4 (Figure 4-13), while the RF model performance was assessed in terms of RMSE and vp-RMSE in Table 5-2. Figure 5-14 compares the capacitance values experimentally measured in Setup 1 and 2 with the predictions obtained from the GP and RF models for the rGO samples listed in Table 5-3. The dashed line indicates perfect agreement between prediction and experiment, and error bars represent the 95% confidence interval of the GP model. For clarity, each sample is assigned a distinct symbol in Table 5-3, and the same symbols are used in Figure 5-14 to represent the corresponding samples.

Table 5-3. Physicochemical characterisation of selected rGO samples.

| Symbol | Temperature (°C) | Time (h) | SSA (m ² /g) | I _D /I _G | O% | Pore volume (cm ³ /g) | Average pore width (nm) |
|--------|------------------|----------|-------------------------|--------------------------------|-------|----------------------------------|-------------------------|
| ● | 80 | 0.5 | 100 | 1.24 | 27.66 | 0.18 | 7.51 |
| ■ | 65 | 2 | 119 | 1.15 | 28.13 | 0.23 | 7.53 |
| ▲ | 72 | 0.8 | 150 | 1.04 | 36.2 | 0.22 | 5.99 |
| ✘ | 72 | 0.8 | 127 | 1.07 | 33.83 | 0.2 | 6.28 |
| + | 80 | 0.5 | 182 | 1.10 | 36 | 0.31 | 6.98 |

Note: N% and S% were negligible in all samples and are therefore not reported in the table.

Notably, the experimentally characterised samples shown in Figure 5-14 lie outside the RF training domain. Consequently, this comparison illustrates model behaviour under extrapolative conditions rather than test-set predictive performance.

As already established in Chapter 4, the GP predictions align closely with the experimental data, with all measurements falling within the model's estimated uncertainty. This behaviour is expected, as the GP was trained directly on experimental data generated under the same synthesis conditions and measurement protocols, and is therefore operating within the domain covered by the training data.

In contrast, the RF model was trained exclusively on literature-derived data from different synthesis routes, characterisation methodologies, and experimental conditions. When applied to the samples synthesised in this work, the RF is consequently forced to extrapolate well beyond its training domain. As a result, the RF consistently overestimates capacitance, yielding only two prediction values: 163.92 F/g for the sample represented by a triangle, and 152.38 F/g for the remaining samples. Relative to

the experimental values, this corresponds to an overestimation ranging from 45.58 F/g (best case, setup 1 for the circle sample) up to 69.29 F/g (worst case, setup 2 for the triangle sample).

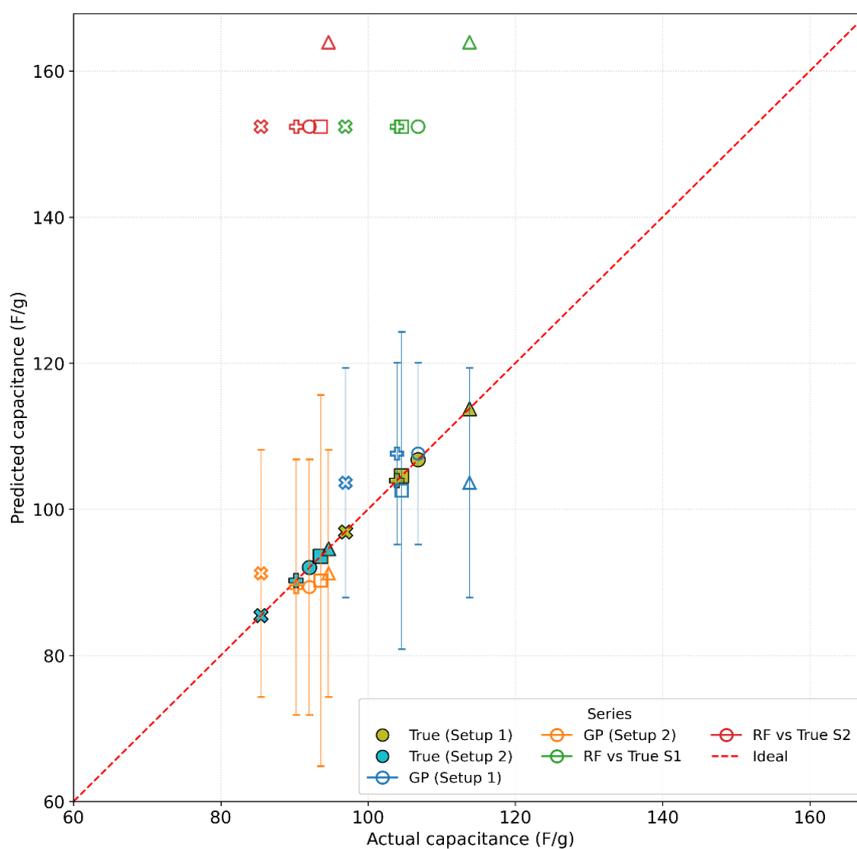


Figure 5-14. Comparison between experimentally measured capacitance values (Setups 1 and 2) and predictions obtained from the Gaussian Process (GP) and Random Forest (RF) models for selected rGO samples described in Table 5-3. Symbols (●, ■, ▲, ✖, +) uniquely identify individual samples listed in Table 5-3 and are used consistently across the table and figure. Error bars on the GP predictions represent the 95% confidence intervals. The dashed line indicates ideal agreement between predicted and measured capacitance.

Several factors contribute to this discrepancy. First, uncertainties inherent to physicochemical characterisation (e.g., spatial variability within a single sample, instrument bias, handling-induced variability during sample preparation) were not quantified in this work but may affect the inputs of the RF. More importantly, these samples are statistical outliers relative to the dataset used for training. The Mahalanobis distance between the data from this work and the literature-based dataset is 4.82 (see Figure 5-15 below), confirming that the RF predictions are obtained under extrapolative conditions, which is a known limitation of tree-based methods [491].

Consequently, this comparison does not constitute a meaningful test-set evaluation of the RF in terms of R^2 but rather illustrates its behaviour beyond the range of the training data. This may explain the systematic overestimation observed. Nevertheless, the magnitude of error remains consistent with the performance metrics previously

established for the RF (RMSE 64.75 ± 24.52 F/g, Table 5-2), underscoring the importance of robust error quantification, a central theme of this chapter.

Although uncertainty estimates can be derived for Random Forest models (e.g., from ensemble variance), they were not included here, as the RF operates in a pronounced out-of-distribution regime for the experimental samples considered. Under such conditions, RF-based uncertainty is known to be poorly calibrated and not directly comparable to the probabilistic uncertainty provided by Gaussian Process models [348,492]. The integration of uncertainty-aware RF approaches is therefore deferred to future work.

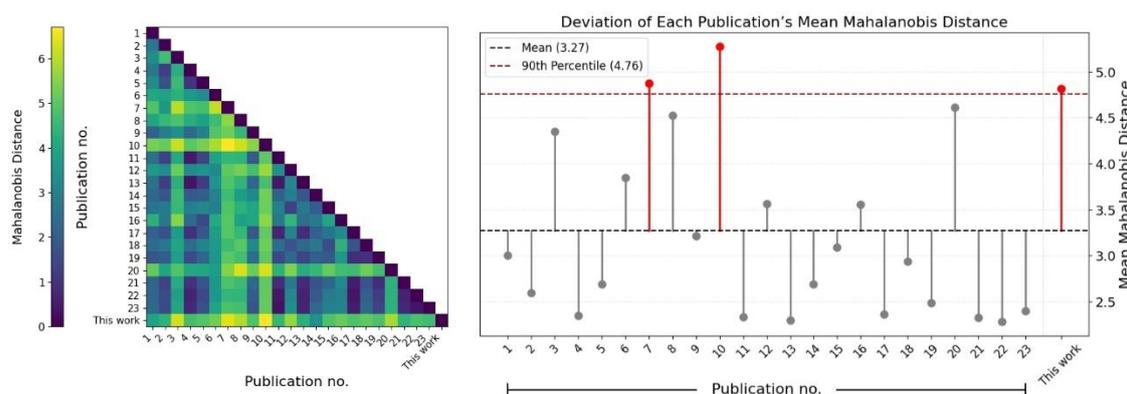


Figure 5-15. Left: Mahalanobis distances across publications from which datapoints were extracted to build the dataset, together with the datapoints from Table 5-3 corresponding to the experimental work described in Chapter 4. Right: Deviation of each publication's mean Mahalanobis distance.

Despite its limited accuracy, the RF provides valuable qualitative insights. In contrast to the GP, it correctly identifies the superior performance of the triangular sample, which also exhibited the highest capacitance in both experimental setups. This sample combines the highest oxygen content with one of the smallest pore diameters and a relatively high surface area. This feature profile aligns with the tendencies highlighted by the SHAP analysis (Figure 5-4) as beneficial for capacitance.

To complement the physicochemical data, SEM imaging was used to gain morphological insights that may help explain the observed capacitance trends. Figure 5-16 presents representative SEM micrographs of the circle and triangle samples. The circle sample appears to form a compact and highly agglomerated structure, where the reduced graphene oxide sheets seem densely restacked and few interlaminar voids are visible. This morphology would be expected to limit accessible surface area and ion diffusion. In contrast, the triangle sample exhibits a more open and wrinkled sheet-like structure with interconnected voids, which might provide increased edge exposure and pathways for electrolyte penetration.

Taken together, the complementary roles of the two models can be summarised as follows: the GP offers robust and uncertainty-aware predictions of capacitance as a function of synthesis parameters within the constraints of the pilot-scale conditions, while the RF provides a broader, more interpretable mapping between physicochemical features and performance, albeit with reduced precision. In combination, these approaches yield a coherent framework for guiding the design of rGO-based supercapacitor electrodes in an industrially scalable context.

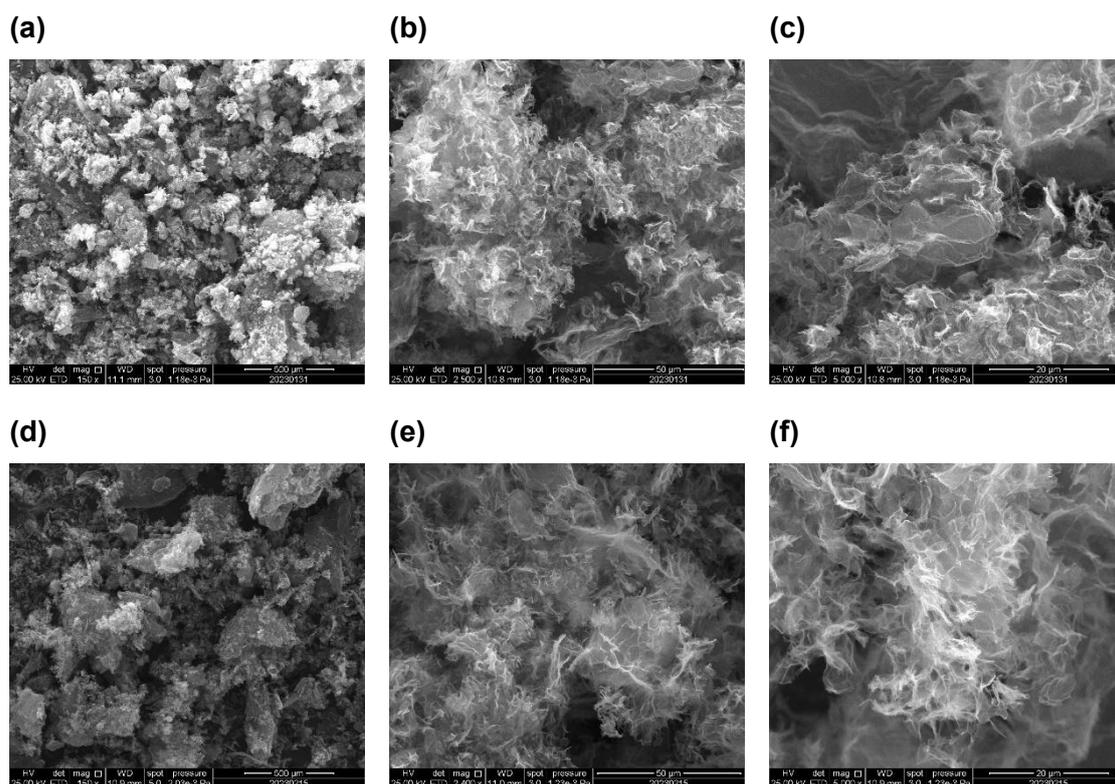


Figure 5-16. Representative SEM micrographs of the circle and triangle samples at three magnifications. (a-c) Circle sample. (d-f) Triangle sample.

5.4 Conclusions and Future Work

This study addresses the complex interplay between physicochemical properties and electrochemical capacitance in reduced graphene oxide (rGO). The findings underscore the context-dependent contribution of key features like oxygen content, surface area, and pore structure to the capacitance of rGO samples.

The differential character of this ML model-building approach lies in the meticulous steps to ensure physical correctness and avoid overfitting. This was achieved by confining the model definition strictly to graphene and its derivatives using KOH and GCD measurements. The data preprocessing stage involved eliminating instances lacking key features like SSA or I_D/I_G ratio and accurately imputing missing values through methods

like Akima interpolation for pore distribution data or zero-filling solely for compositional features.

Moreover, overfitting was avoided through rigorous outlier definition using the Mahalanobis distance for paper outlier identification, Leave-One-Paper-Out Cross Validation (CV) and utilising vp-RMSE during training. Additionally, the assessment of physical correctness was accomplished by aligning the model's conclusions with insights gleaned from papers not included in the datasets, further validating the model's reliability and the general conclusions observed. A comparative analysis with experimental samples synthesised in this work further reinforced these insights. The GP model delivered uncertainty-aware predictions closely matching experimental capacitance, while the RF highlighted structure-performance trends but systematically overestimated values due to extrapolation limitations. Together, they illustrate the trade-off between predictive precision and interpretability.

In summary, the results highlight the key influence of nitrogen, SSA, and oxygen on the capacitance of reduced graphene oxide (rGO). Nitrogen content emerged as the most important factor, aligning with previous studies, and demonstrating a strong impact on capacitance. SSA was shown to have a positive monotonic relationship with capacitance, while average pore width exhibited a context-dependent influence, with an optimal range identified between 2.5 nm and 5 nm. Oxygen content's contribution was found to be variable, influenced by interactions with other features. The analysis also confirmed the importance of maintaining a hierarchically distributed porous structure for effective ion transport. Additionally, the presence of sulphur, although less influential, was positively correlated with enhanced pseudocapacitance and electrical conductivity. Overall, these findings align with previous studies and validate the robustness of the methodologies, reinforcing the reliability of the model presented here.

This study opens several promising research directions to enhance the predictive modelling of rGO supercapacitors. First, expanding the target properties to include volumetric capacitance and cycle-life retention would better address real-world device constraints where electrode density and durability are as critical as specific capacitance. Second, extending the modelling framework to alternative electrolytes, such as organic solvents or ionic liquids, could uncover new design principles for non-aqueous systems while testing the universality of the identified pore-size optima.

To strengthen real-world relevance, future work could integrate experimental validation with model predictions to establish guidelines for material design, such as optimising pore-size distributions or surface functionalisation. The methodology could also be

adapted to directly predict final material properties from synthesis conditions by training surrogate models that map controllable synthesis parameters (e.g., temperature, time, activating agent ratio, electrolyte and measurement setup) directly to electrochemical performance. In such a framework, physicochemical descriptors would be treated as latent variables implicitly learned by the model, with uncertainty-aware approaches such as Gaussian Processes providing confidence bounds to identify regimes where additional characterisation or data acquisition is required. This approach is particularly valuable for rapid material screening and process optimisation, provided sufficient data is available. Additionally, addressing current limitation in the model through larger datasets (to reduce variability) and feature engineering (to test linearity assumptions or identify hidden descriptors) would improve model robustness. Future work should also build on the comparative GP-RF evaluation presented here, expanding the dual-model framework to integrate uncertainty quantification with interpretable feature mapping. Incorporating this hybrid approach in future experimental campaigns could provide both reliable prediction and design guidance, thereby accelerating scalable rGO supercapacitor electrodes. Finally, the methodology's transferability to other electrochemical systems such as catalyst for ORR or CO₂ capture could be explored, leveraging shared principles in structure-property relationships.

6 Conclusions and Future Work

6.1 Conclusions

This thesis demonstrates that, in data-limited experimental settings, accelerating the development of graphene-based materials for electrochemical energy applications requires not only the use of machine learning models, but the ability to embed physical understanding, uncertainty quantification, and experimental constraints directly into the modelling framework. Using reduced graphene oxide as a model system and capacitance as performance metric, the work combines sustainable material choices such as the use of industrial byproducts and non-toxic reducing agents, with physically informed machine learning to enable reproducible and data-efficient optimisation.

A key conclusion is that under data scarcity, experimental noise becomes a dominant factor limiting the reliability of data-driven optimisation. Electrochemical techniques were shown to differ significantly in their sensitivity to experimental variability, with high-uncertainty measurements severely constraining the effectiveness of machine learning approaches when datasets are small. By identifying and selecting robust performance metrics, this thesis established a reliable experimental foundation for probabilistic modelling in low-data regimes.

Building on this foundation, a physics-informed Gaussian Process framework was developed in which experimentally measured noise and physically meaningful constraints were explicitly incorporated. When combined with Bayesian Optimisation, this approach enabled efficient exploration of scalability-constrained synthesis parameter spaces, while maintaining predictive accuracy across multiple laboratories through a physics-informed kernel and noise model that explicitly captured inter-laboratory variability

Although graphene-based materials served as the primary case study, the main contribution of this work is methodological rather than material-specific. The proposed framework is transferable to other electrochemical and materials synthesis systems where datasets are small, experimental variability is high, and scalability and sustainability are key considerations. By reducing experimental waste, incorporating environmentally benign synthesis routes, and providing uncertainty-aware guidance, this work establishes a reproducible and physically grounded machine learning methodology that supports sustainable materials discovery for energy-related applications.

The following sections summarise the main conclusions drawn from Chapters 3 to 5, covering experimental uncertainty analysis, Bayesian optimisation of reduced graphene

oxide capacitance using a physics-informed Gaussian Process framework, and the development of physically informed machine learning models under data scarcity.

6.1.1 Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide

Electrochemical characterisation of graphene-based materials highlights the critical importance of understanding and quantifying experimental uncertainty. In oxygen reduction reaction (ORR) studies conducted using rotating disk electrode (RDE) techniques, considerable variability was observed, with coefficients of variation (CoV) ranging from 1 percent to 129 percent. This variability was largely attributed to subtle inconsistencies in film deposition, ink formulation, and drying conditions, which impacted reproducibility even under controlled protocols. Such high sensitivity to minor changes limits the practicality of RDE for optimisation studies without extensive replication.

In contrast, capacitance measurements using galvanostatic charge-discharge (GCD) testing showed much lower variability, with CoV values generally between 1 and 7 percent. This robustness against sample-to-sample differences, even when variability was intentionally introduced, suggests that GCD testing offers more stable and reproducible performance metrics.

This lower level of uncertainty was a key factor in selecting GCD-based capacitance testing as the foundation for subsequent machine learning-guided optimisation. Given the constrained experimental budget, techniques with higher uncertainty such as RDE would require a much larger number of tests to draw meaningful conclusions. Therefore, the relatively consistent behaviour of GCD measurements provided a practical and statistically justifiable basis for integrating data-driven approaches in material optimisation workflows.

6.1.2 Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes

A physics-informed Gaussian Process model was developed to guide the synthesis of reduced graphene oxide with improved capacitance, using ascorbic acid as a non-toxic and scalable reducing agent. The model incorporated experimentally determined noise to account for variability introduced during reduction, electrode fabrication, and electrochemical testing. Bayesian Optimisation was used to explore the synthesis space efficiently, reducing the number of experimental trials required.

Leave-One-Out Cross-Validation confirmed strong predictive performance, with all capacitance values falling within a 95 percent confidence interval. The observed capacitance remained relatively stable between 100 and 110 F/g over the tested

synthesis window, indicating that time and temperature alone have a weak influence on capacitance within the selected range. Electrochemical measurements using Nyquist plots and cyclic voltammetry supported the pseudocapacitive nature of the material.

To improve generalizability, a second dataset collected under different experimental conditions was integrated by extending the noise model and kernel structure to capture systematic setup-specific variations. The Gaussian Process maintained accuracy despite this added complexity. These results demonstrate the value of combining ML with domain-specific physical insights to enhance reproducibility and control in materials synthesis, offering a data-efficient and scalable strategy for optimizing the performance of electrode materials for energy storage applications.

6.1.3 Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity

The development of a machine learning model for predicting the capacitance of reduced graphene oxide (rGO) has demonstrated that careful model design can bridge the gap between statistical performance and physical relevance. By focusing exclusively on graphene-based materials characterized with KOH electrolyte and galvanostatic charge-discharge methods, the study achieved domain-specific consistency and minimised structural variability. Data preparation was carried out with high attention to scientific integrity, including targeted imputation techniques and the exclusion of entries missing critical features such as surface area or structural disorder indicators.

A distinct emphasis was placed on avoiding overfitting through Leave-One-Paper-Out cross-validation and the use of Mahalanobis distances for identifying literature-based outliers. The training process was further refined using a custom metric that penalized variability in prediction error across references, enhancing both the stability and generalizability of the model. Importantly, model predictions were validated against external literature not included in the training dataset, confirming their physical plausibility.

Key findings highlight nitrogen content, surface area, and pore structure as the dominant features influencing capacitance. Nitrogen doping showed the strongest positive effect, while specific surface area exhibited a clear monotonic relationship. Pore size distribution revealed a performance optimum between 2.5 and 5 nm, reinforcing the role of hierarchical porosity in facilitating effective ion transport and charge storage. A comparative evaluation with experimental samples further highlighted the complementary roles of the models developed in this thesis: the GP offered reliable, uncertainty-aware predictions in line with measured capacitances, while the RF provided

interpretable feature-performance mappings despite systematic overestimation. This dual perspective demonstrate how combining models can balance predictive accuracy with physical insight.

6.2 Future Work

6.2.1 Electrochemical Characterisation and Associated Uncertainty in Kish and Reduced Graphene Oxide

The high variability observed in ORR experiments suggests that further investigation is needed into the physical and procedural factors that influence electrode performance. Quantitative morphological characterisation using scanning electron microscopy (SEM), atomic force microscopy (AFM), or surface profilometry should be employed to better understand how differences in film structure and uniformity correlate with electrochemical behaviour. Such analysis would provide clearer insight into how subtle changes during fabrication impact performance.

Automating key steps in electrode preparation, such as ink deposition and drying, could lead to more consistent results by minimising operator influence. Robotic systems or programmable dispensers offer the potential for standardizing processes across laboratories, improving reproducibility, and enabling more reliable benchmarking of new materials.

From a data analysis perspective, incorporating statistical models that treat performance as a distribution rather than a fixed value would lead to more robust conclusions. Bayesian methods, in particular, are well suited to managing noisy experimental data and can help inform optimisation strategies that are tolerant of uncertainty.

Given the observed reliability of GCD techniques, future studies should consider adopting GCD as a standard method for evaluating graphene-based supercapacitor materials. Further research can expand on this by testing a wider range of carbon-based materials and synthesis conditions under similarly controlled variability.

6.2.2 Bayesian Optimisation of Reduced Graphene Oxide Capacitance using Grey-Box Gaussian Processes

Future efforts could aim to improve feature selection, given that temperature and time had limited influence on capacitance within the tested conditions. Introducing additional synthesis parameters could improve the model's sensitivity and predictive capability. Examples include investigating post-reduction heat treatments, varying the initial oxygen content of graphene oxide, or applying innovative post-processing techniques.

A more detailed analysis of experimental uncertainty is also recommended. Specifically, the role of current density in contributing to measurement noise should be studied through controlled experiments with replicates, which would help quantify intra-condition variability and improve confidence in model predictions. Stability testing under prolonged electrochemical cycling could offer further insight into how synthesis conditions affect long-term performance, which is critical for real-world deployment.

An important and natural extension of this work is the adoption of multi-objective Bayesian optimisation (MOBO). While this thesis focused on maximising capacitance as a single objective, practical materials development requires balancing competing criteria such as capacitance, cycling stability, synthesis time, energy consumption, and production cost. The probabilistic nature of Gaussian Processes makes them particularly well suited for MOBO, as predictive uncertainty can be propagated to construct Pareto fronts rather than single optimal points. Within this context, the existing physics-informed GP framework could be extended to jointly model multiple objectives, enabling informed trade-offs between performance, durability, and costs.

Finally, the methodology developed here can also be applied to other materials systems beyond electrochemical energy storage, including battery cathodes, functional coatings, and bio-based polymers. Broadening the application of this uncertainty-aware, physics-informed machine learning framework could contribute to more reproducible and scalable material development strategies across various scientific and industrial domains.

6.2.3 Modelling Graphene Capacitance with Physics-Informed ML under Data Scarcity

Future research could extend the modelling framework by targeting practical device parameters beyond specific capacitance, such as volumetric energy density and cycle-life retention. These metrics are critical for real-world supercapacitor applications, where compactness and durability are equally important. Additionally, exploring the model's applicability to different electrolyte systems, including organic solvents and ionic liquids, would test the robustness of current conclusions and potentially reveal new pore structure-performance relationships.

A promising direction involves coupling the model with experimental workflows to validate predictions through targeted synthesis and characterisation. Such coupling could be formalised within an automated or semi-automated closed-loop framework, where model predictions inform experimental design and newly generated data are used to iteratively update the model. This approach aligns with emerging concepts in

self-driving laboratories and could significantly reduce experimental cost under data-scarce conditions. Within this context, the model could be used not only for prediction but also for optimisation, for example by identifying synthesis pathways or structural parameter ranges that maximise capacitance while satisfying constraints on stability or manufacturability. Such integration could guide material design more directly, for example by optimising surface functionalisation or tuning pore morphology.

Another avenue lies in shifting the focus from predicting properties based on existing characterisation data to predicting properties directly from synthesis conditions, thereby accelerating high-throughput screening efforts. When combined with Bayesian optimisation or active learning strategies, this shift would enable the automated exploration of synthesis parameter spaces, prioritising experiments with the highest expected information gain or performance improvement.

Improving model robustness will also require addressing current data limitations. Increasing dataset size, introducing replicates, and enhancing feature engineering methods could help test the validity of current assumptions, such as linearity, and may reveal previously overlooked descriptors. Future work could also expand the dual GP-RF framework introduced here, using GP for precise and uncertainty-aware prediction while leveraging RF for interpretable feature analysis. Embedding this hybrid approach into automated experimental validation workflows may accelerate both understanding and optimisation of scalable materials for energy storage applications.

Beyond Gaussian Processes, other ML approaches suitable for data-scarce regimes could be explored. These include Bayesian neural networks and kernel-based methods such as kernel ridge regression with physically motivated kernels, which can incorporate prior knowledge while providing uncertainty estimates. In parallel, additional sources of physical information could be integrated into the modelling framework, such as electrolyte-dependent ion size and solvation effects, descriptors of pore accessibility, electrical conductivity or graphitisation degree, and more detailed surface chemistry metrics derived from advanced characterisation techniques. Incorporating such physically informed descriptors, either explicitly as features or implicitly through model priors and constraints, could further enhance model robustness, interpretability, and transferability to practical device design.

Finally, the modelling methodology could be adapted to related domains such as electrocatalysis or carbon-based CO₂ capture, where structure-property relationships share common patterns and where similar challenges in data quality and physical consistency are present. In these domains, the combination of physically informed ML

with automated optimisation and experiment selection could play a similarly transformative role for accelerating both catalyst and functional material design under sparse and noisy data regimes.

APPENDIX A

This appendix presents additional electrochemical data referenced in Chapter 3. The materials included here support the discussion on experimental variability and reproducibility but are not essential to the main narrative. Each subsection is self-contained and provides context for the corresponding figure or table.

A.1 Variability in RDE Polarisation Curves for Pt40%/Vulcan

This subsection presents RDE polarisation curves obtained using Pt40%/Vulcan catalysts prepared under different deposition conditions to illustrate the resulting variability in electrochemical performance. The data are referenced in Chapter 3 to illustrate the sensitivity of electrochemical parameters to variations in catalyst film preparation, including ink deposition, electrode coverage, and handling.

The results demonstrate that even small differences in experimental execution can lead to measurable variability in electrochemical performance, highlighting the practical motivation for the replication framework discussed in the main text.

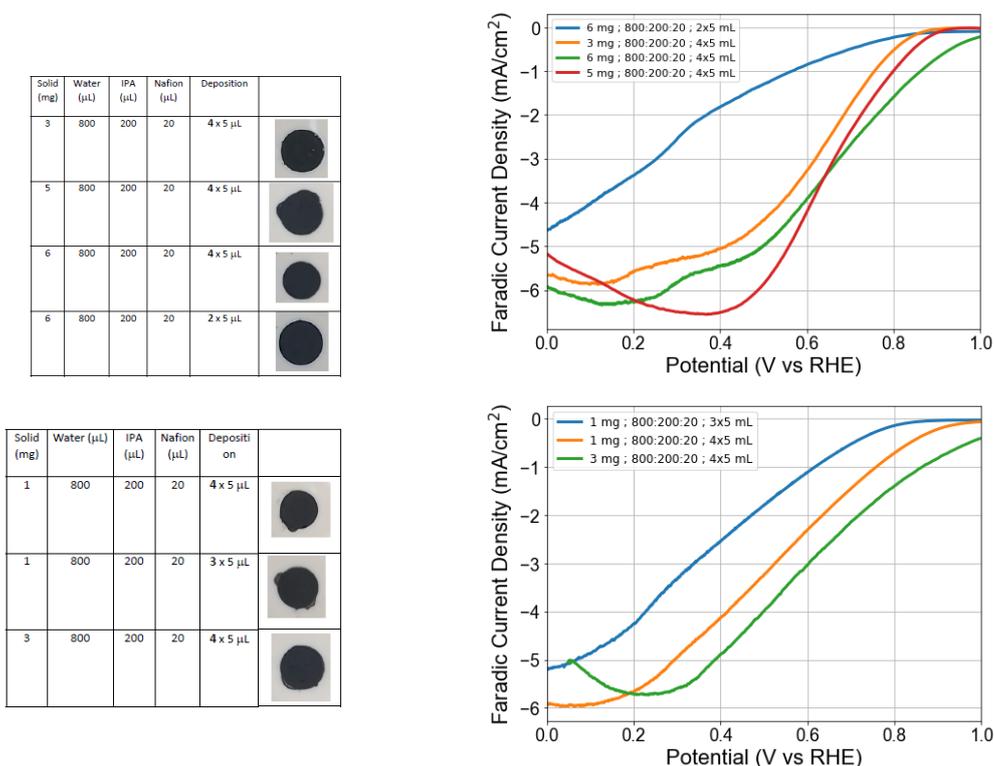


Figure A1. Polarisation curves for using Pt40%/Vulcan different deposition conditions, showing high variability in the results.

A.2 Discharge Time Variability at 1 A/g and 10 A/g Under Different Reduction and Stirring Conditions

This subsection provides galvanostatic discharge (GCD) data for electrodes reduced under varying thermal and stirring conditions, focusing on current densities of 1 A/g and 10 A/g. These results complement the 2 A/g data presented in Chapter 3 (Table 3-6) and are discussed in the context of reproducibility and uncertainty analysis.

Discharge times were measured for electrodes prepared using either magnetic stirring or paddle stirring during the reduction step. The coefficient of variation (CoV, %) was calculated to quantify reproducibility across stirring methods, as defined in Equation 3-2. The data illustrate how reproducibility depends on a combination of current density, reduction temperature, reduction time, and stirring dynamics.

Table A1. Discharge times at 1 A/g and 10 A/g for electrodes reduced under different thermal and stirring conditions. Coefficients of variation (CV, %) indicate reproducibility across stirring methods.

| | Reduction conditions | Stirring | t_0 (s) | $t_{0.4}$ (s) | $t_{0.8}$ (s) | $t_{1.2}$ (s) |
|---------------|-----------------------------|-----------------|-----------|---------------|---------------|---------------|
| 1 A/g | 65°C, 3h | Magnet | 224.44 | 240.03 | 280.69 | 359.32 |
| | 65°C, 3h | Paddle | 182.93 | 205.5 | 252.32 | 331.96 |
| | <i>Coef. Variation (%)</i> | | 14 | 11 | 8 | 6 |
| | 80°C, 0.5h | Magnet | 189.81 | 211.78 | 259.2 | 347.27 |
| | 80°C, 0.5h | Paddle | 184.91 | 202.36 | 244.07 | 329.3 |
| | <i>Coef. Variation (%)</i> | | 2 | 3 | 4 | 4 |
| | 95°C, 0.5h | Magnet | 187.34 | 210.43 | 260.7 | 343.69 |
| | 95°C, 0.5h | Paddle | 200.76 | 223.74 | 274.47 | 359.11 |
| | <i>Coef. Variation (%)</i> | | 5 | 4 | 4 | 3 |
| 10 A/g | 65°C, 3h | Magnet | 10.83 | 11.83 | 14.87 | 20.71 |
| | 65°C, 3h | Paddle | 9.74 | 10.69 | 13.74 | 19.36 |
| | <i>Coef. Variation (%)</i> | | 7 | 7 | 6 | 5 |
| | 80°C, 0.5h | Magnet | 9.41 | 10.41 | 13.46 | 18.67 |
| | 80°C, 0.5h | Paddle | 9.04 | 10.16 | 13.11 | 17.88 |
| | <i>Coef. Variation (%)</i> | | 3 | 2 | 2 | 3 |
| | 95°C, 0.5h | Magnet | 8.96 | 9.59 | 12.32 | 17.78 |
| | 95°C, 0.5h | Paddle | 9.97 | 10.85 | 13.96 | 19.78 |
| | <i>Coef. Variation (%)</i> | | 8 | 9 | 9 | 8 |

APPENDIX B

This appendix provides an independent estimation of the irreducible noise associated with galvanostatic charge-discharge (GCD) measurements, as referenced in Chapter 4. Irreducible noise refers to sources of variability that cannot be attributed to controlled experimental parameters such as temperature fluctuations or mass differences and instead arise from inherent measurement variability and setup-dependent effects.

The results presented here support the assumptions made in the Gaussian Process (GP) uncertainty framework used in Chapter 4, including the estimation of the setup length-scale parameter l_s introduced in Equation 4-14.

B.1 Experimental Estimation of the Irreducible Noise

Irreducible noise refers to factors that are beyond fluctuations in the temperature and mass differences, and which origin we are not explicitly considering. To do this estimation, two electrodes with similar mass (3.008 mg and 3.079 mg) and that incorporate rGO coming from the same reducing conditions (87°C and 120 mins) and heating setup (heating blanket + Arduino) were prepared and characterized. The results are shown in Figure B1 and Table B1 below.

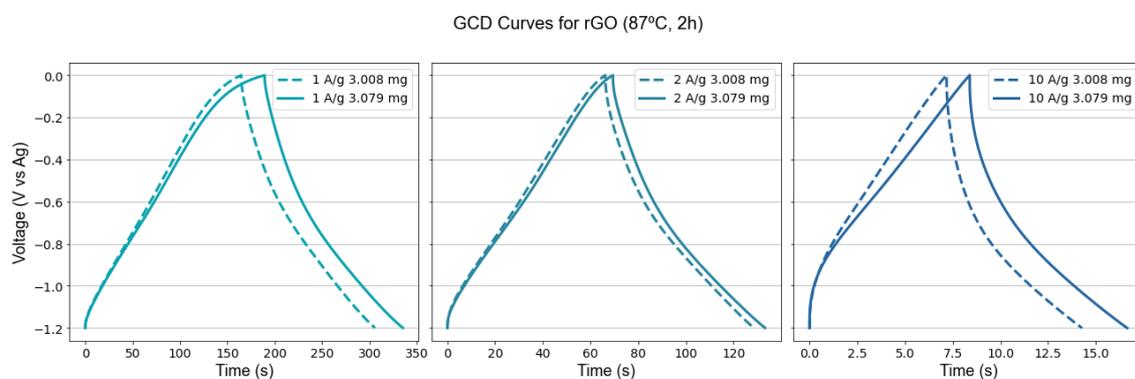


Figure B1 GCD curves for reduced graphene oxide electrodes. The graphene oxide was reduced at 87°C for 120 mins. The curves represent a galvanostatic charge-discharge cycle at 1 A/g (left), 2 A/g (middle) and 10 A/g (right).

Table B1. Calculation of the irreducible noise for 1 A/g, 2 A/g and 10 A/g.

| Current Density Electrode | 1 A/g | 2 A/g | 10 A/g |
|---------------------------------|-------------------|-------------------|------------------|
| 1 | 117.72 F/g | 103.97 F/g | 59.21 F/g |
| 2 | 121.67 F/g | 106.42 F/g | 68.82 F/g |
| $\sqrt{\sigma_{irreducible}^2}$ | 2.79 F/g | 1.73 F/g | 6.80 F/g |

B.2 Derivation of $l_s = 2.3$ from an Assumption of an Inter-Setup Tolerance $\leq 15\%$

This inter-setup tolerance is the result of an estimation considering two contributors to the variance:

- Voltage window difference: which may play an important role, given the pseudocapacitive nature of the rGO materials. A broader voltage window can activate higher number of redox reactions, causing an increase in the specific capacitance measured on the same material. To estimate what is the impact this may cause, the capacitance model from [417] is used. Using the Rule 2 from the Regression Tree described in that model, and calculating the difference in the capacitance obtained for the same material with SSA: $0.1 \text{ m}^2 \cdot \text{mg}^{-1}$, I_D/I_G : 1.07, N%: 0%, O%: 25.7% at PW: 0.8V and 1.2V, the relative difference in the capacitances obtained is 10%.
- Errors associated with data analysis: which may be caused by the approximation of the GCD curves to that of an ideal EDLC for the calculation of the capacitance. In this interlaboratory study [363], capacitances from the same material measured by different groups presented a standard deviation of 5% with respect to the mean value (Fig 5b from [363]).

The expected difference Δy between two setups arises from the Kernel's covariance function:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_y^2 \cdot \exp\left(-\frac{\|\mathbf{x}_i^{[s]} - \mathbf{x}_j^{[s]}\|^2}{2l_s^2}\right) \quad \text{Equation B1}$$

For one-hot encoded setups (where $x_i^{[s]} = [1,0]$ and $x_j^{[s]} = [0,1]$):

- The square distance is $\|\mathbf{x}_i^{[s]} - \mathbf{x}_j^{[s]}\|^2 = 2$
- The covariance between setups becomes:

$$k(\text{Setup 1, Setup 2}) = \sigma_y^2 \cdot \exp\left(-\frac{1}{l_s^2}\right) \quad \text{Equation B2}$$

The expected squared difference between setups is:

$$\begin{aligned} \mathbb{E}[(y_i - y_j)^2] &= k(\mathbf{x}_i, \mathbf{x}_i) + k(\mathbf{x}_j, \mathbf{x}_j) - 2k(\mathbf{x}_i, \mathbf{x}_j) \\ &= 2\sigma_y^2 \left(1 - \exp\left(-\frac{1}{l_s^2}\right)\right) \end{aligned} \quad \text{Equation B3}$$

Assuming zero mean, the expected absolute difference Δy is approximated by the standard deviation:

$$\Delta y \approx \sqrt{\mathbb{E}[(y_i - y_j)^2]} = \sigma_y \sqrt{2 \left(1 - \exp\left(-\frac{1}{l_s^2}\right)\right)} \quad \text{Equation B4}$$

Approximating σ_y to $50 F/g$, based on the capacitance values obtained, Equation B4 gives $l_s = 2.3$.

APPENDIX C

This appendix provides additional methodological details related to data preprocessing and feature interpretation used in the construction and analysis of the dataset presented in Chapter 5. Specifically, it describes the procedure used to estimate missing average pore diameter values, clarifies the role of publication cross-referencing, and provides supplementary visualisations of input feature distributions.

C.1 Estimation of Average Pore Diameter from Published Pore Size Distributions

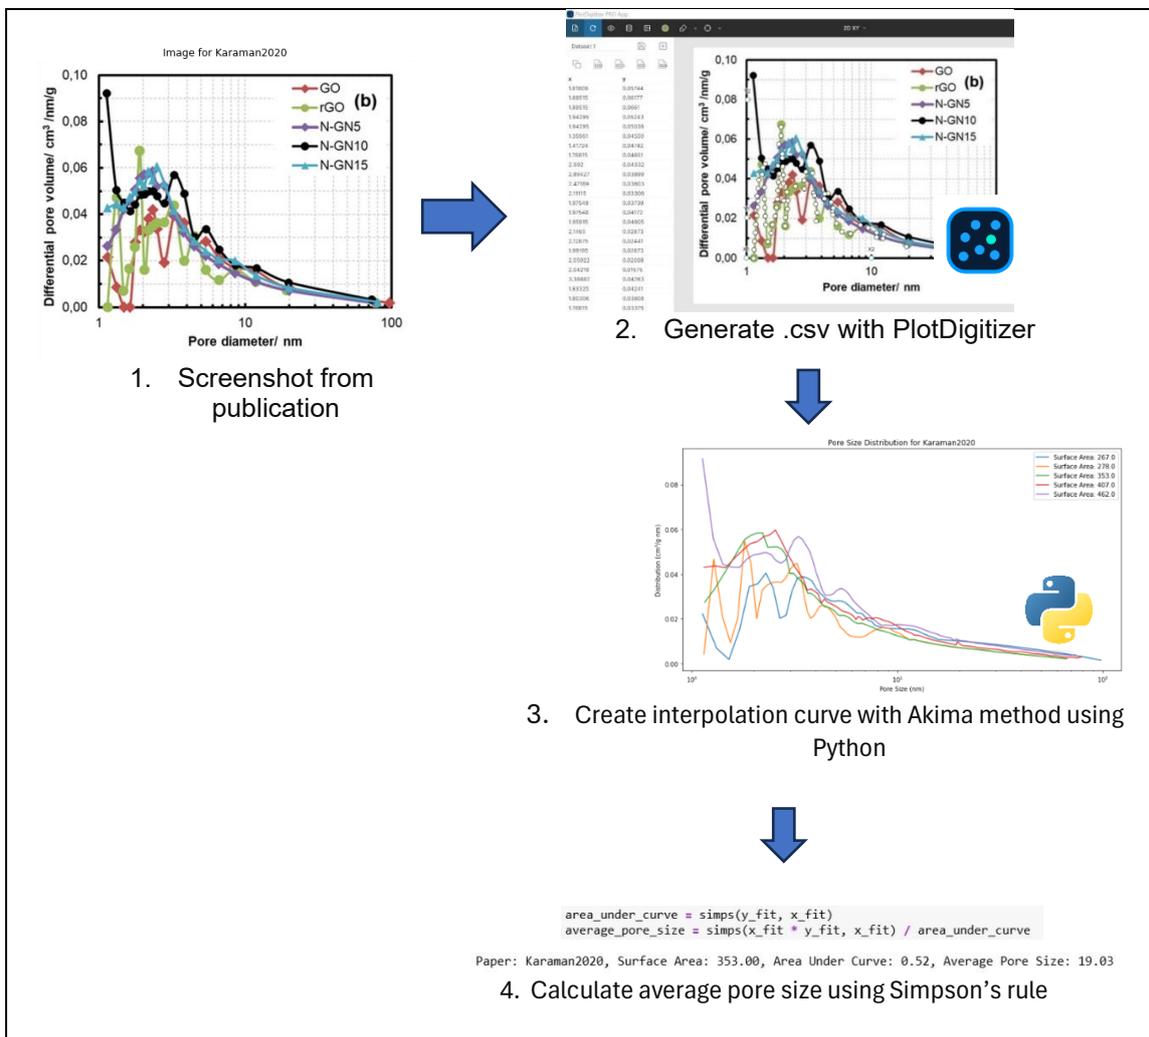


Figure C1. Diagram showing the procedure followed to estimate the average pore size from the available pore size distributions in the referenced publications.

Most of the missing values in the compiled dataset corresponded to the average pore diameter. When pore size distribution curves were available in the original publications, an estimation of this feature was carried out using the procedure illustrated in Figure C1.

The process consisted of the following steps:

1. Digitisation of the pore size distribution curve from the original publication.
2. Interpolation of the extracted data points using Akima interpolation [465] to obtain a smooth approximation of the distribution.
3. Numerical integration of the interpolated curve using Simpson's rule to calculate the area-weighted average pore diameter.

This approach was also used to replace pore size values that had originally been estimated using the empirical relationship derived from Wheeler's theory (Equation 5-1), which assumes a homogeneous distribution of cylindrical pores. As this assumption can lead to a reduction in variance and an oversimplified representation of pore structure [466], the distribution-based estimation method was preferred whenever sufficient data were available.

C.2 Publication Cross-Referencing for Dataset Analysis

To facilitate interpretation of figures that reference multiple publications using numerical identifiers, Table C1 provides a mapping between the publication numbers used throughout the figures in this work and their corresponding full references in the bibliography.

This table is intended to be used in conjunction with figures such as Figure 5-1, which presents Mahalanobis distances between publications, allowing the reader to easily identify the original data sources associated with each indexed entry.

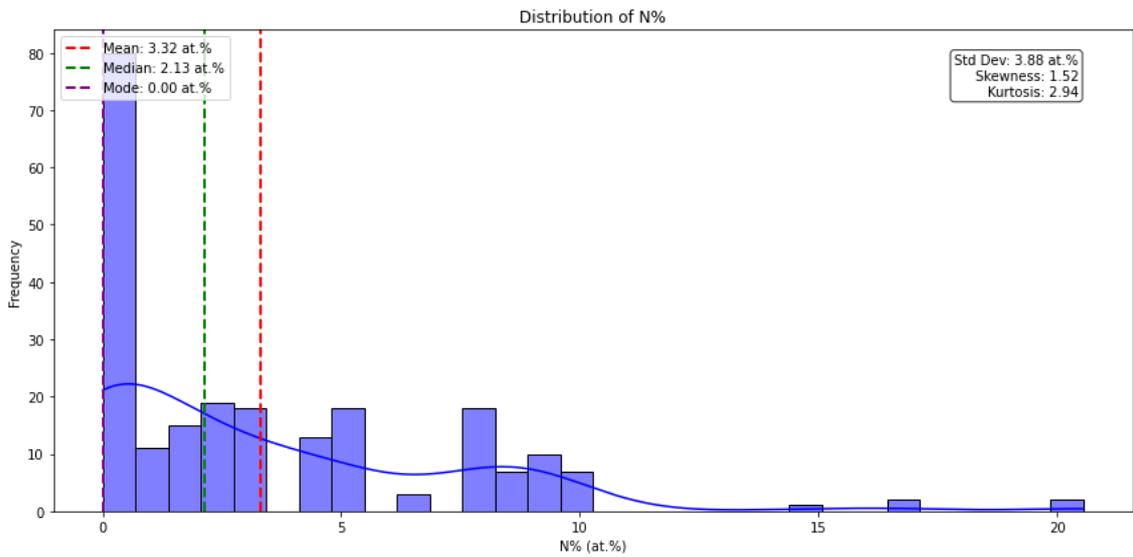
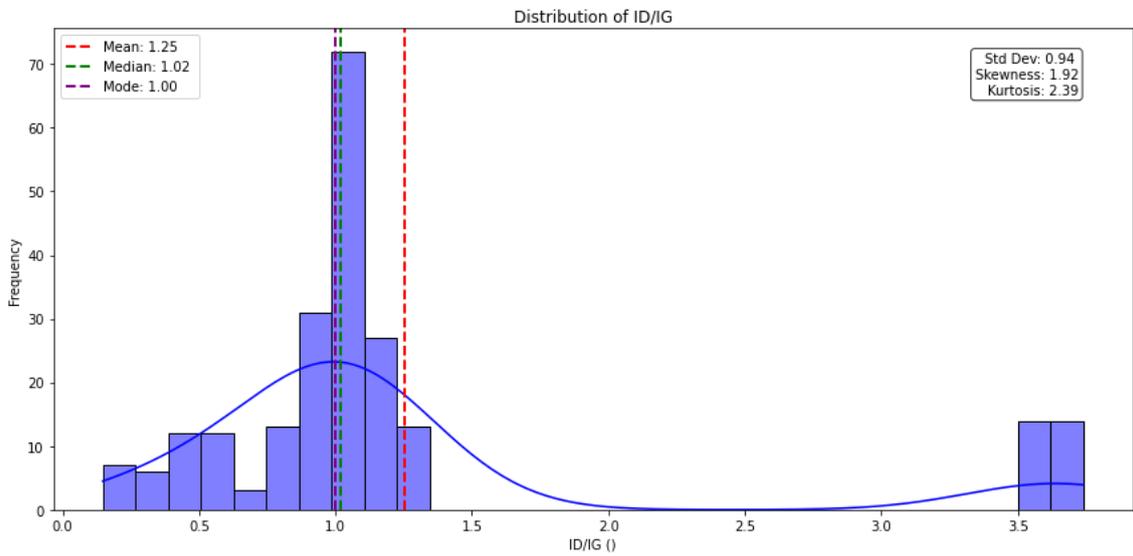
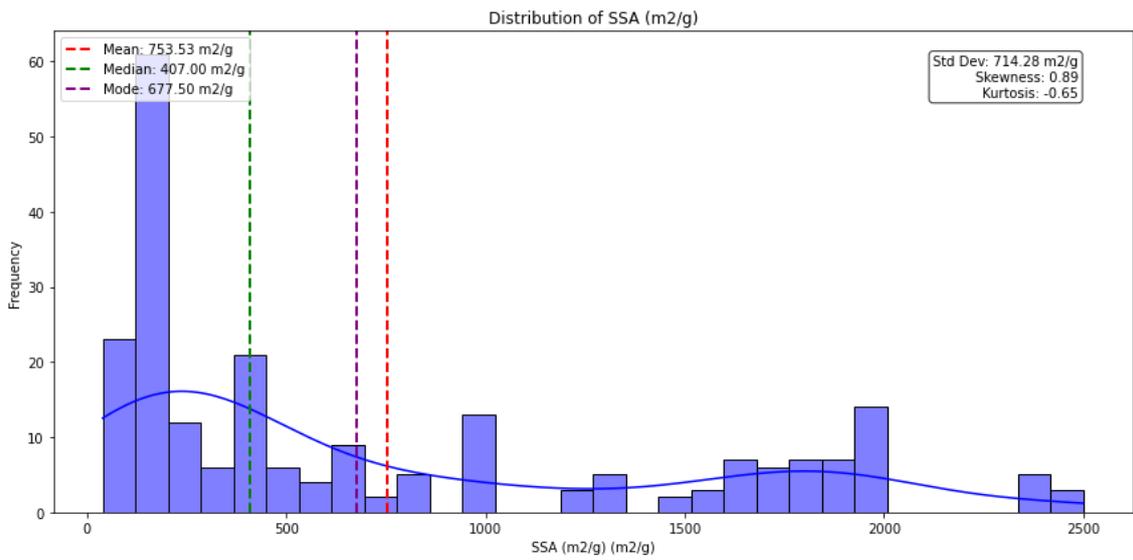
Table C1. On the left: publication numbers shown in the figures of the present work; on the right: their corresponding references in the bibliography.

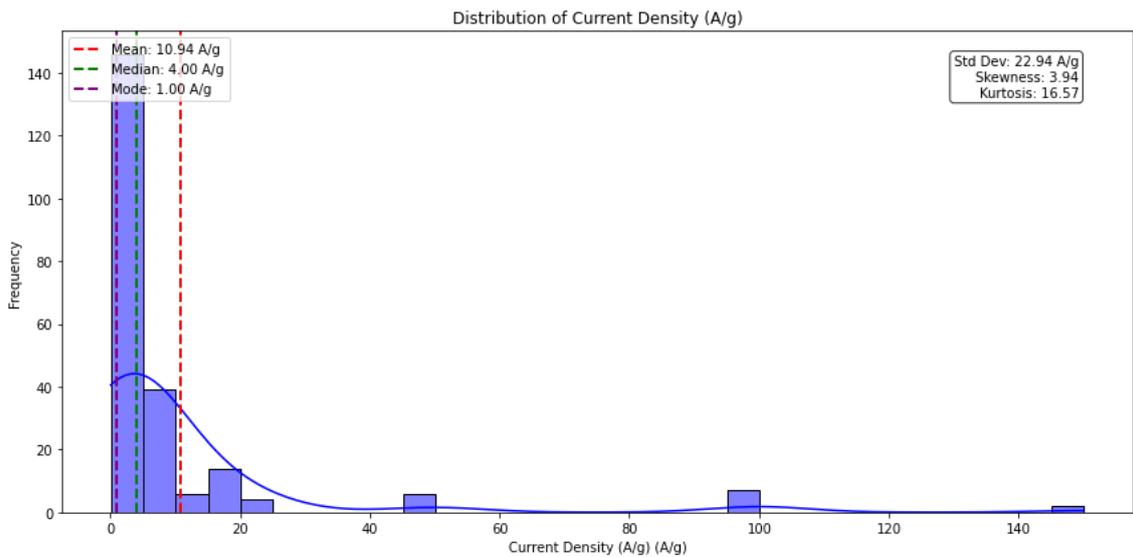
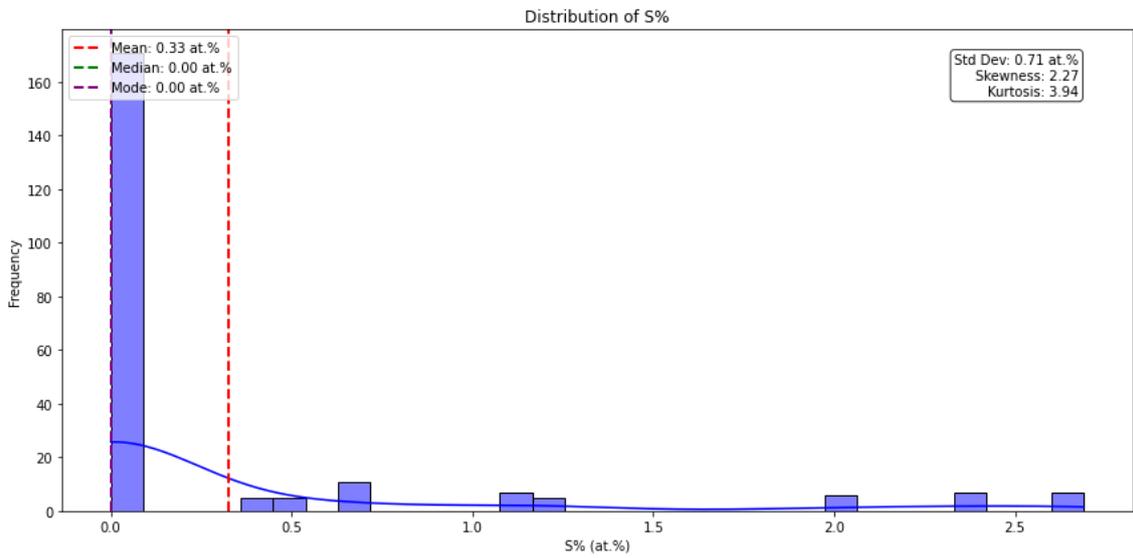
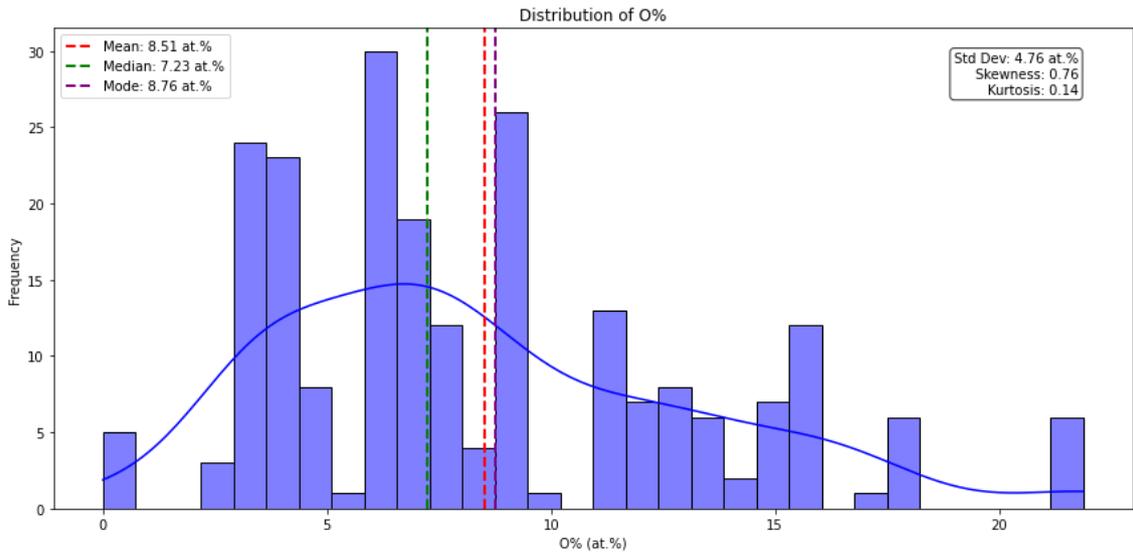
| | |
|----|-------|
| 1 | [493] |
| 2 | [494] |
| 3 | [495] |
| 4 | [496] |
| 5 | [497] |
| 6 | [498] |
| 7 | [499] |
| 8 | [500] |
| 9 | [501] |
| 10 | [502] |
| 11 | [503] |
| 12 | [504] |
| 13 | [505] |
| 14 | [506] |
| 15 | [507] |
| 16 | [508] |
| 17 | [509] |
| 18 | [510] |
| 19 | [511] |
| 20 | [512] |
| 21 | [513] |
| 22 | [514] |
| 23 | [515] |

C.3 Distributions of Input Features and Capacitance

Figure C2 presents the distributions of key input features used in the Gaussian Process models, together with the distribution of the measured capacitance values. These plots complement the discussion in Chapter 5 by providing a broader overview of the statistical properties of the dataset. The vertical dashed lines indicate key summary statistics: the mean value is shown as a red dashed line, the median as a green dashed line, and the mode as a purple dashed line. These markers are included to highlight the central tendency of each distribution and to facilitate comparison between different statistical descriptors, particularly in the presence of skewed or multimodal data.

The boxed annotation located in the upper right corner of each plot reports additional distribution metrics, namely the standard deviation, skewness, and kurtosis, which provide complementary information on the spread, asymmetry, and tail behaviour of the distributions, respectively.





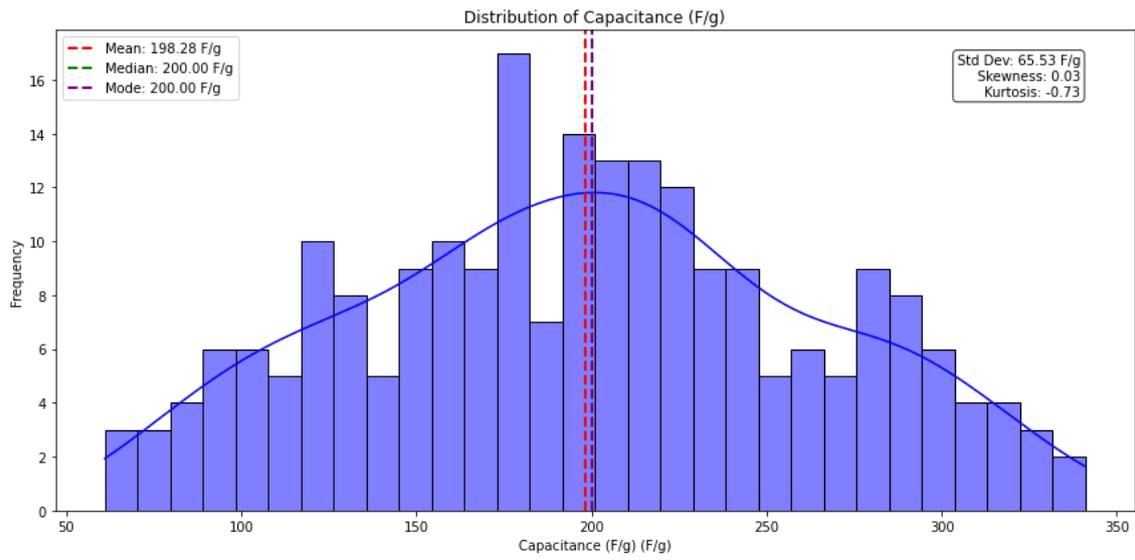


Figure C2. Distributions of input features and Capacitance with key metrics.

REFERENCES

- [1] Council of the European Union. Paris Agreement on climate change. Council of the European Union 2025. <https://www.consilium.europa.eu/en/policies/paris-agreement-climate/> (accessed May 24, 2025).
- [2] European Commission. 2050 long-term strategy. European Commission 2025. https://climate.ec.europa.eu/eu-action/climate-strategies-targets/2050-long-term-strategy_en (accessed May 24, 2025).
- [3] US Department of Energy. DOE Technical Targets for Fuel Cell Systems and Stacks for Transportation Applications | Department of Energy. US Department of Energy 2025. <https://www.energy.gov/eere/fuelcells/doe-technical-targets-fuel-cell-systems-and-stacks-transportation-applications> (accessed May 24, 2025).
- [4] Hanson ED, Mayekar S, Dravid VP. Applying insights from the pharma innovation model to battery commercialization—pros, cons, and pitfalls. *MRS Energy & Sustainability* 2017;4:E10. <https://doi.org/10.1557/MRE.2017.12>.
- [5] Crabtree G. Joint Center for Energy Storage Research JCESR: One Year Later 2014.
- [6] Llewellyn A V., Matruglio A, Brett DJL, Jervis R, Shearing PR. Using In-Situ Laboratory and Synchrotron-Based X-ray Diffraction for Lithium-Ion Batteries Characterization: A Review on Recent Developments. *Condensed Matter* 2020, Vol 5, Page 75 2020;5:75. <https://doi.org/10.3390/CONDMAT5040075>.
- [7] Giannozzi P, Andreussi O, Brumme T, Bunau O, Buongiorno Nardelli M, Calandra M, et al. Advanced capabilities for materials modelling with Quantum ESPRESSO. *Journal of Physics Condensed Matter* 2017;29. <https://doi.org/10.1088/1361-648X/AA8F79>.
- [8] Balasubramani SG, Chen GP, Coriani S, Diedenhofen M, Frank MS, Franzke YJ, et al. TURBOMOLE: Modular program suite for ab initio quantum-chemical and condensed-matter simulations. *Journal of Chemical Physics* 2020;152. <https://doi.org/10.1063/5.0004635>.
- [9] Schmidt J, Marques MRG, Botti S, Marques MAL. Recent advances and applications of machine learning in solid-state materials science. *Npj Computational Materials* 2019 5:1 2019;5:1–36. <https://doi.org/10.1038/s41524-019-0221-0>.
- [10] Peng J, Schwalbe-Koda D, Akkiraju K, Xie T, Giordano L, Yu Y, et al. Human- and machine-centred designs of molecules and materials for sustainability and decarbonization. *Nature Reviews Materials* 2022 7:12 2022;7:991–1009. <https://doi.org/10.1038/s41578-022-00466-5>.
- [11] Elsevier. Reaxys | An expert-curated chemical database. Elsevier 2025. <https://www.elsevier.com/products/reaxys> (accessed May 24, 2025).
- [12] The Materials Project. Lawrence Berkeley National Laboratory 2025. <https://next-gen.materialsproject.org/> (accessed May 24, 2025).
- [13] OQMD. Northwestern University 2025. <https://oqmd.org/> (accessed May 24, 2025).

- [14] NOMAD Laboratory. NOMAD. NOMAD Laboratory 2025. <https://nomad-lab.eu/nomad-lab/> (accessed May 24, 2025).
- [15] Aflow - Automatic FLOW for Materials Discovery. AFLOW Consortium 2025. <https://www.afowlib.org/> (accessed May 24, 2025).
- [16] Kononova O, Huo H, He T, Rong Z, Botari T, Sun W, et al. Text-mined dataset of inorganic materials synthesis recipes. *Sci Data* 2019;6:1–11. <https://doi.org/10.1038/s41597-019-0297-x>.
- [17] Kasilingam S, Yang R, Singh SK, Farahani MA, Rai R, Wuest T. Physics-based and data-driven hybrid modeling in manufacturing: a review. *Prod Manuf Res* 2024;12. <https://doi.org/10.1080/21693277.2024.2305358>.
- [18] ArcelorMittal. Who we are - ArcelorMittal Europe. ArcelorMittal Europe 2025. <https://europe.arcelormittal.com/aboutarcelormittaleurope/who/europewhoweare> (accessed May 24, 2025).
- [19] World Steel Association. About steel. World Steel Association Website 2025. <https://worldsteel.org/about-steel> (accessed September 11, 2022).
- [20] Bauer I, Hans-Joachim K. Iron Catalysis in Organic Chemistry. In: Plietker B, editor. *Iron Catalysis in Organic Chemistry: Reactions and Applications*, Wiley-VCH Verlag; 2008, p. 1–27.
- [21] Silver J. *Chemistry of Iron*. Glasgow, UK: Blackie Academic & Professional; 1993.
- [22] Metallos. Fe Pourbaix diagram. <https://commons.wikimedia.org/wiki/File:Fe-Pourbaix-DiagramSv> 2007.
- [23] Ao X, Zhang W, Li Z, Li JG, Soule L, Huang X, et al. Markedly enhanced oxygen reduction activity of single-atom Fe catalysts via integration with Fe nanoclusters. *ACS Nano* 2019;13:11853–62. <https://doi.org/10.1021/ACS.NANO.9B05913>.
- [24] Jiang WJ, Gu L, Li L, Zhang Y, Zhang X, Zhang LJ, et al. Understanding the High Activity of Fe-N-C Electrocatalysts in Oxygen Reduction: Fe/Fe₃C Nanoparticles Boost the Activity of Fe-N_x. *J Am Chem Soc* 2016;138:3570–8. <https://doi.org/10.1021/JACS.6B00757>.
- [25] Na Y, Wei P, Zhou L. Photochemical Hydrogen Generation Initiated by Oxidative Quenching of the Excited Ru(bpy)₃ (2+) * by a Bio-Inspired [2Fe₂S] Complex. *Chemistry* 2016;22:10365–8. <https://doi.org/10.1002/CHEM.201600541>.
- [26] Lennox AJJ, Fischer S, Jurrat M, Luo SP, Rockstroh N, Junge H, et al. Copper-Based Photosensitisers in Water Reduction: A More Efficient In Situ Formed System and Improved Mechanistic Understanding. *Chemistry – A European Journal* 2016;22:1233–8. <https://doi.org/10.1002/CHEM.201503812>.
- [27] Goy R, Bertini L, Rudolph T, Lin S, Schulz M, Zampella G, et al. Photocatalytic Hydrogen Evolution Driven by [FeFe] Hydrogenase Models Tethered to Fluorene and Silafluorene Sensitizers. *Chemistry – A European Journal* 2017;23:334–45. <https://doi.org/10.1002/CHEM.201603140>.
- [28] Nurttala SS, Becker R, Hessels J, Woutersen S, Reek JNH. Photocatalytic Hydrogen Evolution by a Synthetic [FeFe] Hydrogenase Mimic Encapsulated in a Porphyrin Cage. *Chemistry – A European Journal* 2018;24:16395–406. <https://doi.org/10.1002/CHEM.201803351>.

- [29] Mun J, Lee MJ, Park JW, Oh DJ, Lee DY, Doo SG. Non-aqueous redox flow batteries with nickel and iron tris(2,2'-bipyridine) complex electrolyte. *Electrochemical and Solid-State Letters* 2012;15:A80. <https://doi.org/10.1149/2.033206ESL/XML>.
- [30] Bamgbopa MO, Shao-Horn Y, Almheiri S. The potential of non-aqueous redox flow batteries as fast-charging capable energy storage solutions: demonstration with an iron–chromium acetylacetonate chemistry. *J Mater Chem A Mater* 2017;5:13457–68. <https://doi.org/10.1039/C7TA02022H>.
- [31] Xing X, Zhao Y, Li Y. A non-aqueous redox flow battery based on tris(1,10-phenanthroline) complexes of iron(II) and cobalt(II). *J Power Sources* 2015;293:778–83. <https://doi.org/10.1016/J.JPOWSOUR.2015.06.016>.
- [32] Shin M, Noh C, Chung Y, Kwon Y. All iron aqueous redox flow batteries using organometallic complexes consisting of iron and 3-[bis (2-hydroxyethyl)amino]-2-hydroxypropanesulfonic acid ligand and ferrocyanide as redox couple. *Chemical Engineering Journal* 2020;398:125631. <https://doi.org/10.1016/J.CEJ.2020.125631>.
- [33] Sharma S, Andrade GA, Maurya S, Popov IA, Batista ER, Davis BL, et al. Iron-iminopyridine complexes as charge carriers for non-aqueous redox flow battery applications. *Energy Storage Mater* 2021;37:576–86. <https://doi.org/10.1016/J.ENSM.2021.01.035>.
- [34] Zwicker MFR, Moghadam M, Zhang W, Nielsen C V. Automotive battery pack manufacturing – a review of battery to tab joining. *Journal of Advanced Joining Processes* 2020;1:100017. <https://doi.org/10.1016/J.JAJP.2020.100017>.
- [35] Ramesh P, Lenin NC. High Power Density Electrical Machines for Electric Vehicles-Comprehensive Review Based on Material Technology. *IEEE Trans Magn* 2019;55. <https://doi.org/10.1109/TMAG.2019.2929145>.
- [36] ArcelorMittal. Electrical steels - Industry. ArcelorMittal Corporate Website 2022. <https://industry.arcelormittal.com/electricalsteels> (accessed September 16, 2022).
- [37] ArcelorMittal. Energy. ArcelorMittal Corporate Website 2022. <https://corporate.arcelormittal.com/industries/energy> (accessed September 16, 2022).
- [38] Borit NV. Fuel cells & electrolyzers. Borit NV 2022. <https://www.borit.be/offerings/fuel-cell-electrolyzers> (accessed September 16, 2022).
- [39] Encyclopaedia Britannica. Iron processing – The metal. Encyclopaedia Britannica, Inc 2022. <https://www.britannica.com/technology/iron-processing/The-metal> (accessed September 17, 2022).
- [40] Liu S, Loper CR. Kish, a source of crystalline graphite. *Carbon N Y* 1991;29:1119–24. [https://doi.org/10.1016/0008-6223\(91\)90029-I](https://doi.org/10.1016/0008-6223(91)90029-I).
- [41] Liu S, Loper CR. The formation of kish graphite. *Carbon N Y* 1991;29:547–55. [https://doi.org/10.1016/0008-6223\(91\)90119-4](https://doi.org/10.1016/0008-6223(91)90119-4).
- [42] Frost R. The recovery of kish graphite from secondary sources. University of Birmingham, 2015.

- [43] The Iron Carbon Phase Diagram n.d. https://www.tf.uni-kiel.de/matwis/amat/iss/kap_6/illustr/s6_1_2.html (accessed September 17, 2022).
- [44] Novoselov KS, Geim AK, Morozov S V., Jiang D, Zhang Y, Dubonos S V., et al. Electric field in atomically thin carbon films. *Science* (1979) 2004;306:666–9. <https://doi.org/10.1126/SCIENCE.1102896>.
- [45] Science of Fullerenes and Carbon Nanotubes. *Science of Fullerenes and Carbon Nanotubes* 1996. <https://doi.org/10.1016/B978-0-12-221820-0.X5000-X>.
- [46] Wolf EL. *Graphene: A New Paradigm in Condensed Matter and Device Physics* 2013. <https://doi.org/10.1093/ACPROF:OSO/9780199645862.001.0001>.
- [47] Song HS, Li SL, Miyazaki H, Sato S, Hayashi K, Yamada A, et al. Origin of the relatively low transport mobility of graphene grown through chemical vapor deposition. *Scientific Reports* 2012 2:1 2012;2:1–6. <https://doi.org/10.1038/srep00337>.
- [48] An JC, Kim HJ, Hong I. Preparation of Kish graphite-based graphene nanoplatelets by GIC (graphite intercalation compound) via process. *Journal of Industrial and Engineering Chemistry* 2015;26:55–60. <https://doi.org/10.1016/J.JIEC.2014.12.016>.
- [49] Wang S, Kravchyk K V., Krumeich F, Kovalenko M V. Kish Graphite Flakes as a Cathode Material for an Aluminum Chloride-Graphite Battery. *ACS Appl Mater Interfaces* 2017;9:28478–85. <https://doi.org/10.1021/ACSAMI.7B07499>.
- [50] Technology | Form Energy n.d. <https://formenergy.com/technology/> (accessed September 16, 2022).
- [51] ESS Iron Flow Chemistry | ESS, Inc. n.d. <https://essinc.com/iron-flow-chemistry/> (accessed September 16, 2022).
- [52] Fan J, Chen M, Zhao Z, Zhang Z, Ye S, Xu S, et al. Bridging the gap between highly active oxygen reduction reaction catalysts and effective catalyst layers for proton exchange membrane fuel cells. *Nat Energy* 2021;6. <https://doi.org/10.1038/s41560-021-00824-7>.
- [53] James BD. DOE Hydrogen and Fuel Cells Program Review Presentation. 2019.
- [54] Papageorgopoulos D. Fuel Cell R&D Overview. 2019.
- [55] Li C, Zhao DH, Long HL, Li M. Recent advances in carbonized non-noble metal–organic frameworks for electrochemical catalyst of oxygen reduction reaction. *Rare Metals* 2021;40. <https://doi.org/10.1007/s12598-020-01694-w>.
- [56] Osmieri L, Park J, Cullen DA, Zelenay P, Myers DJ, Neyerlin KC. Status and challenges for the application of platinum group metal-free catalysts in proton-exchange membrane fuel cells. *Curr Opin Electrochem* 2021;25. <https://doi.org/10.1016/j.coelec.2020.08.009>.
- [57] Li P, Wang HL. Recent advances in carbon-supported iron group electrocatalysts for the oxygen reduction reaction. *Xinxing Tan Cailiao/New Carbon Materials* 2021;36. [https://doi.org/10.1016/S1872-5805\(21\)60072-0](https://doi.org/10.1016/S1872-5805(21)60072-0).

- [58] Jin J, Yin J, Liu H, Lu M, Li J, Tian M, et al. Transition Metal (Fe, Co and Ni)–Carbide–Nitride (M–C–N) Nanocatalysts: Structure and Electrocatalytic Applications. *ChemCatChem* 2019;11. <https://doi.org/10.1002/cctc.201900570>.
- [59] Wang Y, Li J, Wei Z. Transition-metal-oxide-based catalysts for the oxygen reduction reaction. *J Mater Chem A Mater* 2018;6. <https://doi.org/10.1039/c8ta01321g>.
- [60] Xu Z, Zhao H, Liang J, Wang Y, Li T, Luo Y, et al. Noble-metal-free electrospun nanomaterials as electrocatalysts for oxygen reduction reaction. *Materials Today Physics* 2020;15. <https://doi.org/10.1016/j.mtphys.2020.100280>.
- [61] Li R, Li X, Liu Q, Yue H, Meng Y, Tian Y. Effect of Fe–N–C single atom/cluster catalyst on ORR of PEMFC. *Int J Hydrogen Energy* 2025;139:718–29. <https://doi.org/10.1016/j.ijhydene.2025.05.294>.
- [62] Li Y, Chen B, Duan X, Chen S, Liu D, Zang K, et al. Atomically dispersed Fe-N-P-C complex electrocatalysts for superior oxygen reduction. *Appl Catal B* 2019;249. <https://doi.org/10.1016/j.apcatb.2019.03.016>.
- [63] Domínguez C, Pérez-Alonso FJ, Abdel Salam M, Gómez De La Fuente JL, Al-Thabaiti SA, Basahel SN, et al. Effect of transition metal (M: Fe, Co or Mn) for the oxygen reduction reaction with non-precious metal catalysts in acid medium. *Int J Hydrogen Energy*, vol. 39, 2014, p. 5309–18. <https://doi.org/10.1016/j.ijhydene.2013.12.078>.
- [64] Luo Y, Du X, Wu L, Wang Y, Li J, Ricardez-Sandoval L. Machine-Learning-Accelerated Screening of Double-Atom/Cluster Electrocatalysts for the Oxygen Reduction Reaction. *Journal of Physical Chemistry C* 2023;127:20372–84. <https://doi.org/10.1021/acs.jpcc.3c05753>.
- [65] Wu G, Zelenay P. Activity versus stability of atomically dispersed transition-metal electrocatalysts. *Nat Rev Mater* 2024;9:643–56. <https://doi.org/10.1038/s41578-024-00703-z>.
- [66] Xu X, Zhang X, Kuang Z, Xia Z, Rykov AI, Yu S, et al. Investigation on the demetallation of Fe-N-C for oxygen reduction reaction: The influence of structure and structural evolution of active site. *Appl Catal B* 2022;309:121290. <https://doi.org/https://doi.org/10.1016/j.apcatb.2022.121290>.
- [67] Holby EF, Wang G, Zelenay P. Acid Stability and Demetalation of PGM-Free ORR Electrocatalyst Structures from Density Functional Theory: A Model for “Single-Atom Catalyst” Dissolution. *ACS Catal* 2020;10:14527–39. <https://doi.org/10.1021/acscatal.0c02856>.
- [68] Chi K, Wang Z, Sun T, He P, Xiao F, Lu J, et al. Simultaneously Engineering the First and Second Coordination Shells of Single Iron Catalysts for Enhanced Oxygen Reduction. *Small* 2024. <https://doi.org/10.1002/smll.202311817>.
- [69] Sun Y, Silvioli L, Sahraie NR, Ju W, Li J, Zitolo A, et al. Activity-Selectivity Trends in the Electrochemical Production of Hydrogen Peroxide over Single-Site Metal-Nitrogen-Carbon Catalysts. *J Am Chem Soc* 2019;141:12372–81. <https://doi.org/10.1021/jacs.9b05576>.

- [70] Shukla AK, Hariprakash B. Chemistry, Electrochemistry, and Electrochemical Applications | Iron. Encyclopedia of Electrochemical Power Sources, 2009. <https://doi.org/10.1016/B978-044452745-5.00063-0>.
- [71] Guo L, Zhang Y, Wang J, Ma L, Ma S, Zhang Y, et al. Unlocking the energy capabilities of micron-sized LiFePO₄. Nature Communications 2015 6:1 2015;6:1–8. <https://doi.org/10.1038/ncomms8898>.
- [72] Lin J, Sun Y-H, Lin X. Metal-organic framework-derived LiFePO₄ cathode encapsulated in O,F-codoped carbon matrix towards superior lithium storage. Nano Energy 2022;91:106655. <https://doi.org/https://doi.org/10.1016/j.nanoen.2021.106655>.
- [73] Wang B, Xie Y, Liu T, Luo H, Wang B, Wang C, et al. LiFePO₄ quantum-dots composite synthesized by a general microreactor strategy for ultra-high-rate lithium ion batteries. Nano Energy 2017;42:363–72. <https://doi.org/https://doi.org/10.1016/j.nanoen.2017.11.040>.
- [74] Zhang Q, Zhou Y, Tong Y, Chi Y, Liu R, Dai C, et al. Reduced Graphene Oxide Coating LiFePO₄ Composite Cathodes for Advanced Lithium-Ion Battery Applications. Int J Mol Sci 2023;24. <https://doi.org/10.3390/ijms242417549>.
- [75] Cai Y, Huang D, Ma Z, Wang H, Huang Y, Wu X, et al. Construction of highly conductive network for improving electrochemical performance of lithium iron phosphate. Electrochim Acta 2019;305:563–70. <https://doi.org/10.1016/J.ELECTACTA.2019.02.114>.
- [76] Du G, Xi Y, Tian X, Zhu Y, Zhou Y, Deng C, et al. One-step hydrothermal synthesis of 3D porous microspherical LiFePO₄/graphene aerogel composite for lithium-ion batteries. Ceram Int 2019;45:18247–54. <https://doi.org/10.1016/J.CERAMINT.2019.06.035>.
- [77] Chien WC, Li YR, Wu SH, Wu YS, Wu ZH, Li YJJ, et al. Modifying the morphology and structure of graphene oxide provides high-performance LiFePO₄/C/rGO composite cathode materials. Advanced Powder Technology 2020;31:4541–51. <https://doi.org/10.1016/J.APT.2020.10.002>.
- [78] Ma Z, Zuo Z, Li L, Li Y, Ma Z, Li L, et al. Unleash the Capacity Potential of LiFePO₄ through Rocking-Chair Coordination Chemistry. Adv Funct Mater 2021;2108692. <https://doi.org/10.1002/ADFM.202108692>.
- [79] Zhao Q, Zhang S, Wang C, Jiang G, Hu M, Chen Z, et al. Combined First-principles and Experimental Studies of a V-doped LiFePO₄/C Composite as a Cathode Material for Lithium-ion Batteries. Int J Electrochem Sci 2021;16:21053. <https://doi.org/10.20964/2021.05.01>.
- [80] Robinson JB, Xi K, Kumar RV, Ferrari AC, Au H, Titirici M-M, et al. 2021 roadmap on lithium sulfur batteries. Journal of Physics: Energy 2021;3:31501. <https://doi.org/10.1088/2515-7655/abdb9a>.
- [81] Jin X, Gao S, Wu A, Zhao J, Huang H, Cao G. Dual-constrained sulfur in FeS₂@C nanostructured lithium-sulfide batteries. ACS Appl Energy Mater 2020;3:10950–60. https://doi.org/10.1021/ACSAEM.0C01929/SUPPL_FILE/AE0C01929_SI_001.PDF.

- [82] Xi Y, Ye X, Duan S, Li T, Zhang J, Jia L, et al. Iron vacancies and surface modulation of iron disulfide nanoflowers as a high power/energy density cathode for ultralong-life stable Li storage. *J Mater Chem A Mater* 2020;8:14769–77. <https://doi.org/10.1039/D0TA04038J>.
- [83] Dewald GF, Liaqat Z, Lange MA, Tremel W, Zeier WG. Influence of Iron Sulfide Nanoparticle Sizes in Solid-State Batteries**. *Angewandte Chemie* 2021;133:18096–100. <https://doi.org/10.1002/ANGE.202106018>.
- [84] Shen Y, Zhang Q, Wang Y, Gu L, Zhao X, Shen X, et al. A Pyrite Iron Disulfide Cathode with a Copper Current Collector for High-Energy Reversible Magnesium-Ion Storage. *Advanced Materials* 2021;33:2103881. <https://doi.org/10.1002/ADMA.202103881>.
- [85] Wu H, Xia G, Yu X. Unlocking the Potential of Iron Sulfides for Sodium-Ion Batteries by Ultrafine Pulverization. *Small* 2024. <https://doi.org/10.1002/smll.202312190>.
- [86] Wang YX, Yang J, Chou SL, Liu HK, Zhang WX, Zhao D, et al. Uniform yolk-shell iron sulfide-carbon nanospheres for superior sodium-iron sulfide batteries. *Nat Commun* 2015;6. <https://doi.org/10.1038/ncomms9689>.
- [87] Soloveichik GL. Flow Batteries: Current Status and Trends. *Chem Rev* 2015;115:11533–58. <https://doi.org/10.1021/cr500720t>.
- [88] Ponce de León C, Frías-Ferrer A, González-García J, Szánto DA, Walsh FC. Redox flow cells for energy conversion. *J Power Sources* 2006;160:716–32. <https://doi.org/10.1016/J.JPOWSOUR.2006.02.095>.
- [89] Sánchez-Díez E, Ventosa E, Guarnieri M, Trovò A, Flox C, Marcilla R, et al. Redox flow batteries: Status and perspective towards sustainable stationary energy storage. *J Power Sources* 2021;481:228804. <https://doi.org/10.1016/J.JPOWSOUR.2020.228804>.
- [90] Sun C, Zhang H. Review of the Development of First-Generation Redox Flow Batteries: Iron-Chromium System. *ChemSusChem* 2022;15. <https://doi.org/10.1002/cssc.202101798>.
- [91] Raytheon, EnerVault. Final Technical Report Flow Battery Solution for Smart Grid Applications. 2015.
- [92] Xin Z. New energy-storing tech at forefront of nation's transition. *China Daily* 2023. <https://www.chinadaily.com.cn/a/202304/13/WS64373b1fa31057c47ebb9cbe.html> (accessed January 2, 2026).
- [93] Zeng YK, Zhou XL, Zeng L, Yan XH, Zhao TS. Performance enhancement of iron-chromium redox flow batteries by employing interdigitated flow fields. *J Power Sources* 2016;327:258–64. <https://doi.org/https://doi.org/10.1016/j.jpowsour.2016.07.066>.
- [94] Che H, Gao Y, Yang J, Hong S, Hao L, Xu L, et al. Bismuth nanoparticles anchored on N-doped graphite felts to give stable and efficient iron-chromium redox flow batteries. *New Carbon Materials* 2024;39:131–41. [https://doi.org/https://doi.org/10.1016/S1872-5805\(24\)60837-1](https://doi.org/https://doi.org/10.1016/S1872-5805(24)60837-1).

- [95] Wan CT-C, Rodby KE, Perry ML, Chiang Y-M, Brushett FR. Hydrogen evolution mitigation in iron-chromium redox flow batteries via electrochemical purification of the electrolyte. *J Power Sources* 2023;554:232248. <https://doi.org/10.1016/j.jpowsour.2022.232248>.
- [96] Zeng Y, Zhou X, Jiang H, Ren Y, Zhao T. High-Performance Flow-Field Structured Iron-Chromium Redox Flow Batteries for Large-Scale Energy Storage. *ECS Meeting Abstracts* 2017;MA2017-01:179. <https://doi.org/10.1149/MA2017-01/2/179>.
- [97] Zeng YK, Zhao TS, Zhou XL, Zeng L, Wei L. The effects of design parameters on the charge-discharge performance of iron-chromium redox flow batteries. *Appl Energy* 2016;182:204–9. <https://doi.org/10.1016/J.APENERGY.2016.08.135>.
- [98] Zhang H, Chen N, Sun C, Luo X. Investigations on physicochemical properties and electrochemical performance of graphite felt and carbon felt for iron-chromium redox flow battery. *Int J Energy Res* 2020;44:3839–53. <https://doi.org/10.1002/ER.5179>.
- [99] Chen N, Zhang H, Luo XD, Sun CY. SiO₂-decorated graphite felt electrode by silicic acid etching for iron-chromium redox flow battery. *Electrochim Acta* 2020;336. <https://doi.org/10.1016/J.ELECTACTA.2020.135646>.
- [100] Xu N, Li X, Zhao X, Goodenough JB, Huang K. A novel solid oxide redox flow battery for grid energy storage. *Energy Environ Sci* 2011;4:4942–6. <https://doi.org/10.1039/C1EE02489B>.
- [101] Gong K, Ma X, Conforti KM, Kuttler KJ, Grunewald JB, Yeager KL, et al. A zinc-iron redox-flow battery under \$100 per kW h of system capital cost. *Energy Environ Sci* 2015;8:2941–5. <https://doi.org/10.1039/C5EE02315G>.
- [102] Arroyo-Currás N, Hall JW, Dick JE, Jones RA, Bard AJ. An Alkaline Flow Battery Based on the Coordination Chemistry of Iron and Cobalt. *J Electrochem Soc* 2014;162:A378. <https://doi.org/10.1149/2.0461503JES>.
- [103] Wei X, Xia G-G, Kirby B, Thomsen E, Li B, Nie Z, et al. An Aqueous Redox Flow Battery Based on Neutral Alkali Metal Ferri/ferrocyanide and Polysulfide Electrolytes. *J Electrochem Soc* 2016;163:A5150–3. <https://doi.org/10.1149/2.0221601JES>.
- [104] Zeng YK, Zhao TS, Zhou XL, Wei L, Jiang HR. A low-cost iron-cadmium redox flow battery for large-scale energy storage. *J Power Sources* 2016;330:55–60. <https://doi.org/10.1016/J.JPOWSOUR.2016.08.107>.
- [105] Hruska LW, Savinell RF. Investigation of Factors Affecting Performance of the Iron-Redox Battery. *J Electrochem Soc* 1981;128:18. <https://doi.org/10.1149/1.2127366>.
- [106] Lloyd D, Vainikka T, Ronkainen M, Kontturi K. Characterisation and application of the Fe(II)/Fe(III) redox reaction in an ionic liquid analogue. *Electrochim Acta* 2013;109:843–51. <https://doi.org/10.1016/J.ELECTACTA.2013.08.013>.
- [107] Gong K, Xu F, Grunewald JB, Ma X, Zhao Y, Gu S, et al. All-Soluble All-Iron Aqueous Redox-Flow Battery. *ACS Energy Lett* 2016;1:89–93. <https://doi.org/10.1021/ACSENERGYLETT.6B00049>.

- [108] Zeng YK, Zhou XL, An L, Wei L, Zhao TS. A high-performance flow-field structured iron-chromium redox flow battery. *J Power Sources* 2016;324:738–44. <https://doi.org/10.1016/J.JPOWSOUR.2016.05.138>.
- [109] Wang W, Nie Z, Chen B, Chen F, Luo Q, Wei X, et al. A new Fe/V redox flow battery using a sulfuric/chloric mixed-acid supporting electrolyte. *Adv Energy Mater* 2012;2:487–93. <https://doi.org/10.1002/AENM.201100527>.
- [110] Oei D-G. Chemically regenerative redox fuel cells. *Journal of Applied Electrochemistry* 1982 12:1 1982;12:41–51. <https://doi.org/10.1007/BF01112063>.
- [111] Zeng YK, Zhao TS, An L, Zhou XL, Wei L. A comparative study of all-vanadium and iron-chromium redox flow batteries for large-scale energy storage. *J Power Sources* 2015;300:438–43. <https://doi.org/10.1016/J.JPOWSOUR.2015.09.100>.
- [112] Licht S, Wang B, Ghosh S. Energetic Iron(VI) Chemistry: The Super-Iron Battery. *Science* (1979) 1999;285:1039–42. <https://doi.org/10.1126/science.285.5430.1039>.
- [113] Farmand M, Jiang D, Wang B, Ghosh S, Ramaker DE, Licht S. Super-iron nanoparticles with facile cathodic charge transfer. *Electrochem Commun* 2011;13:909–12. <https://doi.org/https://doi.org/10.1016/j.elecom.2011.03.039>.
- [114] Zhang Y, Zhao X, Zhang S, Zhang G, Liu S. Optimized preparation conditions of yttria doped zirconia coatings on potassium ferrate (VI) electrode for alkaline super-iron battery. *Appl Energy* 2012;99:265–71. <https://doi.org/https://doi.org/10.1016/j.apenergy.2012.05.036>.
- [115] Huang J, Yang Z, Wang S, Feng Z. Preparation and stability study of potassium ferrate (VI) coated with phthalocyanine for alkaline super-iron battery. *Journal of Solid State Electrochemistry* 2015;19:723–30. <https://doi.org/10.1007/s10008-014-2658-x>.
- [116] Wang S, Wang Y, Chen S, Hou H, Li H. Enhanced capacity and stability of K₂FeO₄ cathode with poly(3-hexylthiophene) coating for alkaline super-iron battery. *Electrochim Acta* 2016;213:132–9. <https://doi.org/https://doi.org/10.1016/j.electacta.2016.06.164>.
- [117] Yan C, Zhu L, Dong J, Gu D, Jiang H, Wang B. Structural modification of isomorphous (SO₄)₂-doped K₂FeO₄ for remediating the stability and enhancing the discharge of super-iron battery. *R Soc Open Sci* 2019;6:180919. <https://doi.org/10.1098/rsos.180919>.
- [118] Licht S. A High Capacity Li-Ion Cathode: The Fe(III/VI) Super-Iron Cathode. *Energies (Basel)* 2010;3:960–72. <https://doi.org/10.3390/en3050960>.
- [119] Licht S, Ghosh S, Naschitz V, Halperin N, Halperin L. Fe(VI) Catalyzed Manganese Redox Chemistry: Permanganate and Super-Iron Alkaline Batteries. *J Phys Chem B* 2001;105:11933–6. <https://doi.org/10.1021/jp012178t>.
- [120] Licht S, Naschitz V, Ghosh S. Silver Mediation of Fe(VI) Charge Transfer: Activation of the K₂FeO₄ Super-iron Cathode. *J Phys Chem B* 2002;106:5947–55. <https://doi.org/10.1021/jp014720j>.
- [121] Yu X, Licht S. Advances in Fe(VI) charge storage: Part II. Reversible alkaline super-iron batteries and nonaqueous super-iron batteries. *J Power Sources*

- 2007;171:1010–22.
<https://doi.org/https://doi.org/10.1016/j.jpowsour.2007.06.064>.
- [122] Licht S, Ghosh S. High power BaFe(VI)O₄/MnO₂ composite cathode alkaline super-iron batteries. *J Power Sources* 2002;109:465–8. [https://doi.org/https://doi.org/10.1016/S0378-7753\(02\)00113-1](https://doi.org/https://doi.org/10.1016/S0378-7753(02)00113-1).
- [123] Licht S, Ghosh S, Naschitz V. Hydroxide Activated AgMnO₄ Alkaline Cathodes, Alone and in Combination with Fe(VI) Super-Iron, BaFeO₄. *Electrochemical and Solid-State Letters* 2001;4:A209. <https://doi.org/10.1149/1.1413202>.
- [124] Yu X, Licht S. Zirconia coating stabilized super-iron alkaline cathodes. *J Power Sources* 2007;173:1012–6. <https://doi.org/https://doi.org/10.1016/j.jpowsour.2007.08.041>.
- [125] Licht S, Yu X, Zheng D. Cathodic chemistry of high performance Zr coated alkaline materials. *Chem Commun* 2006:4341–3. <https://doi.org/10.1039/B608716G>.
- [126] Gasgoo. GAC Group launches super iron lithium battery. <https://AutonewsGasgooCom/m/70020712Html> 2022.
- [127] Egashira M. SECONDARY BATTERIES – METAL-AIR SYSTEMS | Iron–Air (Secondary and Primary). *Encyclopedia of Electrochemical Power Sources* 2009:372–5. <https://doi.org/10.1016/B978-044452745-5.00102-7>.
- [128] McKerracher RD, Poncedeleon C, Wills RGA, Shah AA, Walsh FC. A Review of the Iron–Air Secondary Battery for Energy Storage. *Chempluschem* 2015;80:323–35. <https://doi.org/10.1002/CPLU.201402238>.
- [129] Form energy. Form energy - Battery Technology. <https://FormenergyCom/Technology/Battery-Technology/> n.d.
- [130] EnergyTech. DOE Awards \$70M to Xcel Energy for 2 Iron-Air Long Duration Energy Storage Systems. <https://www.energytech.com/energy-storage/article/21274397/doe-awards-70m-to-xcel-energy-for-2-iron-air-long-duration-energy-storage-systems> 2023.
- [131] Jiang J, Liu J. Iron anode-based aqueous electrochemical energy storage devices: Recent advances and future perspectives 2767. <https://doi.org/10.1002/idm2.12007>.
- [132] Shukla AK, Hariprakash B. Secondary batteries – nickel systems | Nickel–Iron. *Encyclopedia of Electrochemical Power Sources* 2009:522–7. <https://doi.org/10.1016/B978-044452745-5.00155-6>.
- [133] Manohar AK, Yang C, Malkhandi S, Prakash GKS, Narayanan SR. Enhancing the Performance of the Rechargeable Iron Electrode in Alkaline Batteries with Bismuth Oxide and Iron Sulfide Additives. *J Electrochem Soc* 2013;160:A2078–84. <https://doi.org/10.1149/2.066311JES>.
- [134] Shukla AK, Hariprakash B. Secondary batteries – nickel systems | Electrodes: Iron. *Encyclopedia of Electrochemical Power Sources* 2009:418–23. <https://doi.org/10.1016/B978-044452745-5.00163-5>.
- [135] Posada JOG, Hall PJ. The Effect of Electrolyte Additives on the Performance of Iron Based Anodes for NiFe Cells. *J Electrochem Soc* 2015;162:A2036–43. <https://doi.org/10.1149/2.0451510JES>.

- [136] Arunkumar PS, Maiyalagan T, Kheawhom S, Mao S, Jiang Z. Effect of carbon material additives on hydrogen evolution at rechargeable alkaline iron battery electrodes. *Mater Sci Energy Technol* 2021;4:236–41. <https://doi.org/10.1016/J.MSET.2021.06.007>.
- [137] Gil Posada JO, Hall PJ. Post-hoc comparisons among iron electrode formulations based on bismuth, bismuth sulphide, iron sulphide, and potassium sulphide under strong alkaline conditions. *J Power Sources* 2014;268:810–5. <https://doi.org/10.1016/J.JPOWSOUR.2014.06.126>.
- [138] Manohar AK, Yang C, Narayanan SR. The Role of Sulfide Additives in Achieving Long Cycle Life Rechargeable Iron Electrodes in Alkaline Batteries. *J Electrochem Soc* 2015;162:A1864–72. <https://doi.org/10.1149/2.0741509JES>.
- [139] Sun Z, Kong L, Wang X, Lei Y, Li X, Tang H, et al. Study of the electrochemical properties of Fe₃O₄/C–Bi composites prepared by spray drying and their use in cylindrical sealed nickel iron batteries. *Int J Hydrogen Energy* 2023;48:9785–96. <https://doi.org/10.1016/J.IJHYDENE.2022.11.241>.
- [140] Abdalla AH, Oseghale CI, GilPosada JO, Hall PJ. Rechargeable nickel–iron batteries for large-scale energy storage. *IET Renewable Power Generation* 2016;10:1529–34. <https://doi.org/10.1049/IET-RPG.2016.0051>.
- [141] Hussain F, Rahman MZ, Sivasengaran AN, Hasanuzzaman M. Energy storage technologies. *Energy for Sustainable Development: Demand, Supply, Conversion and Management* 2020:125–65. <https://doi.org/10.1016/B978-0-12-814645-3.00006-7>.
- [142] Kong D, Wang Y, Huang S, Zhang B, Lim Y Von, Sim GJ, et al. 3D Printed Compressible Quasi-Solid-State Nickel-Iron Battery. *ACS Nano* 2020;14:9675–86. https://doi.org/10.1021/ACSNANO.0C01157/SUPPL_FILE/NN0C01157_SI_001.PDF.
- [143] Barton JP, Gammon RJL, Rahil A. Characterisation of a Nickel-iron Battolyser, an Integrated Battery and Electrolyser. *Front Energy Res* 2020;8:318. <https://doi.org/10.3389/FENRG.2020.509052>.
- [144] ArcelorMittal. S-in motion® Battery Pack for BEV. https://AutomotiveArcelormittalCom/s-In_motion_solutions/Battery_pack 2023.
- [145] Tata Steel. HILUMIN® for cylindrical battery cell casings in electric vehicles 2023. <https://www.tatasteeleurope.com/automotive/key-products/hilumin> (accessed October 22, 2023).
- [146] Takahashi T, Ishizuka K, Kawanishi K. Properties of Nickel-Coated Steel Sheets for Battery Case. 2015.
- [147] Mills CA, Batyrev E, Jansen MJR, Ahmad M, Pathan TS, Legge EJ, et al. Improvement in the Electrical Properties of Nickel-Plated Steel Using Graphitic Carbon Coatings. *Adv Eng Mater* 2019;21. <https://doi.org/10.1002/adem.201900408>.
- [148] Batyrev E. Graphene coating: A long-life time performance to battery steel. *Open Access Journal of Science and Technology* 2020.

- [149] Mehlmann B, Olowinsky A, Thuilot M, Gillner A. Spatially modulated laser beam micro welding of CuSn6 and nickel-plated DC04 steel for battery applications. *Journal of Laser Micro Nanoengineering* 2014;9:276–81. <https://doi.org/10.2961/jlmn.2014.03.0019>.
- [150] Sadeghian A, Iqbal N. A review on dissimilar laser welding of steel-copper, steel-aluminum, aluminum-copper, and steel-nickel for electric vehicle battery manufacturing. *Opt Laser Technol* 2022;146. <https://doi.org/10.1016/j.optlastec.2021.107595>.
- [151] Indhu R, Divya S, Tak M, Soundarapandian S. Microstructure development in Pulsed Laser Welding of Dual Phase Steel to Aluminium Alloy. vol. 26, Elsevier B.V.; 2018, p. 495–502. <https://doi.org/10.1016/j.promfg.2018.07.058>.
- [152] Trinh LN, Lee D. The characteristics of laser welding of a thin aluminum tab and steel battery case for lithium-ion battery. *Metals (Basel)* 2020;10:1–15. <https://doi.org/10.3390/met10060842>.
- [153] European Hydrogen Backbone (EHB) Project 2020. <https://ehb.eu/> (accessed January 2, 2026).
- [154] European Hydrogen Backbone. How a dedicated Hydrogen Infrastructure can be created. 2020.
- [155] Georges C, Sturel T, Drillet P, Mataigne JM. Absorption/desorption of diffusible hydrogen in aluminized boron steel. *ISIJ International* 2013;53:1295–304. <https://doi.org/10.2355/isijinternational.53.1295>.
- [156] Allen QS, Nelson TW. Microstructural evaluation of hydrogen embrittlement and successive recovery in advanced high strength steel. *J Mater Process Technol* 2019;265:12–9. <https://doi.org/10.1016/j.jmatprotec.2018.09.039>.
- [157] Ohaeri EG, Qin W, Szpunar J. A critical perspective on pipeline processing and failure risks in hydrogen service conditions. *J Alloys Compd* 2021;857. <https://doi.org/10.1016/j.jallcom.2020.158240>.
- [158] Yang J, Fan R, Li J, Meng F. Hydrogen Leakage Detectors Based on a Polymer Microfiber Decorated with Pd Nanoparticles. *IEEE Sens J* 2019;19:6736–41. <https://doi.org/10.1109/JSEN.2019.2912867>.
- [159] Cai S, Liu F, Wang R, Xiao Y, Li K, Caucheteur C, et al. Narrow bandwidth fiber-optic spectral combs for renewable hydrogen detection. *Science China Information Sciences* 2020;63. <https://doi.org/10.1007/s11432-020-3058-2>.
- [160] Wang G, Dai J, Yang M. Fiber-Optic Hydrogen Sensors: A Review. *IEEE Sens J* 2021;21:12706–18. <https://doi.org/10.1109/JSEN.2020.3029519>.
- [161] Liu RX, Zhang X, Zhu XS, Shi YW. Optical Fiber Hydrogen Sensor Based on the EVA/Pd Coated Hollow Fiber. *IEEE Photonics J* 2022;14. <https://doi.org/10.1109/JPHOT.2022.3168009>.
- [162] Lee S, Ryu B, Kim I, Song YW. Temperature- and Ambient Pressure-Independent Sensing of Hydrogen in Fluids Using Cascaded Interferometers Incorporated in Optical Fibers. *Adv Mater Technol* 2023;8. <https://doi.org/10.1002/admt.202201273>.

- [163] Martinez JE, Johannes LB, Yayathi S, Figuered JM, Bilc ZM. Metallography of Battery Resistance Spot Welds. n.d.
- [164] Taheri P, Hsieh S, Bahrami M. Investigating electrical contact resistance losses in lithium-ion battery assemblies for hybrid and electric vehicles. *J Power Sources* 2011;196:6525–33.
<https://doi.org/https://doi.org/10.1016/j.jpowsour.2011.03.056>.
- [165] Behera P, Rajagopalan SK, Brahimi S, Venturella CA, Gaydos SP, Straw RJ, et al. Effect of brush plating process variables on the microstructures of Cd and ZnNi coatings and hydrogen embrittlement. *Surf Coat Technol* 2021;417.
<https://doi.org/10.1016/j.surfcoat.2021.127181>.
- [166] Sadananda K, Yang JH, Iyyer N, Phan N, Rahman A. Sacrificial Zn-Ni coatings by electroplating and hydrogen embrittlement of high-strength steels. *Corrosion Reviews* 2021;39:487–517. <https://doi.org/10.1515/corrrev-2021-0038>.
- [167] Wetegrove M, Duarte MJ, Taube K, Rohloff M, Gopalan H, Scheu C, et al. Preventing Hydrogen Embrittlement: The Role of Barrier Coatings for the Hydrogen Economy. *Hydrogen* 2023;4:307–22.
<https://doi.org/10.3390/hydrogen4020022>.
- [168] Shi K, Xiao S, Ruan Q, Wu H, Chen G, Zhou C, et al. Hydrogen permeation behavior and mechanism of multi-layered graphene coatings and mitigation of hydrogen embrittlement of pipe steel. *Appl Surf Sci* 2022;573.
<https://doi.org/10.1016/j.apsusc.2021.151529>.
- [169] Pradhan A, Vishwakarma M, Dwivedi SK. A review: The impact of hydrogen embrittlement on the fatigue strength of high strength steel. *Mater Today Proc*, vol. 26, Elsevier Ltd; 2019, p. 3015–9.
<https://doi.org/10.1016/j.matpr.2020.02.627>.
- [170] Sun B, Wang D, Lu X, Wan D, Ponge D, Zhang X. Current challenges and opportunities towards understanding hydrogen embrittlement mechanisms in advanced high-strength steels. *Acta Metallurgica Sinica Letters* 2021;34:741–54.
- [171] Gong P, Turk A, Nutter J, Yu F, Wynne B, Rivera-Diaz-del-Castillo P, et al. Hydrogen embrittlement mechanisms in advanced high strength steel. *Acta Mater* 2022;223. <https://doi.org/10.1016/j.actamat.2021.117488>.
- [172] Mohrbacher H, Kern A. Nickel Alloying in Carbon Steel: Fundamentals and Applications. *Alloys* 2023;2:1–28.
- [173] Kwon YJ, Lee T, Lee J, Chun YS, Lee CS. Role of Cu on hydrogen embrittlement behavior in Fe–Mn–C–Cu TWIP steel. *Int J Hydrogen Energy* 2015;40:7409–19.
<https://doi.org/https://doi.org/10.1016/j.ijhydene.2015.04.022>.
- [174] Venezuela J, Liu Q, Zhang M, Zhou Q, Atrens A. A review of hydrogen embrittlement of martensitic advanced high-strength steels 2016;34:153–86.
<https://doi.org/doi:10.1515/corrrev-2016-0006>.
- [175] Okonkwo PC, Barhoumi EM, Ben Belgacem I, Mansir IB, Aliyu M, Emori W, et al. A focused review of the hydrogen storage tank embrittlement mechanism process. *Int J Hydrogen Energy* 2023;48:12935–48.
<https://doi.org/https://doi.org/10.1016/j.ijhydene.2022.12.252>.

- [176] Koyama M, Akiyama E, Lee Y-K, Raabe D, Tsuzaki K. Overview of hydrogen embrittlement in high-Mn steels. *Int J Hydrogen Energy* 2017;42:12706–23. <https://doi.org/https://doi.org/10.1016/j.ijhydene.2017.02.214>.
- [177] Unno H, Nagata T, Fujimoto N, Fukuda M. Fe-based Metal Foils for Current Collectors in Li-ion Secondary Batteries. *Nippon Steel Technical Report* 2019;122.
- [178] Tata Steel. HILUMIN® Ultra-clean nickel-plated steel Production Process Pickling Plating Cold Rolling Temper Rolling Annealing Slitting/Packaging. 2013.
- [179] Takahashi T, Matsumura K, Ishizuka K, Goto Y. Properties of Flexible Nickel Coated Steel Sheets for Battery Case. 2015.
- [180] Outokumpu Oyj. Stainless steel stacks up for fuel cell bipolar plates. Outokumpu Oyj 2021. <https://www.outokumpu.com/en/expertise/2021/stainless-steel-stacks-up-for-fuel-cell-bipolar-plates>.
- [181] C.D. Wälzholz GmbH & Co. KG. Steel for bipolar plates. CD Wälzholz GmbH & Co KG 2025. <https://www.waelzholz.com/en/sustainability/technologies/future-steel-materials/steel-for-bipolar-plates>.
- [182] Liu R, Jia Q, Zhang B, Lai Z, Chen L. Protective coatings for metal bipolar plates of fuel cells: A review. *Int J Hydrogen Energy* 2022;47:22915–37. <https://doi.org/10.1016/j.ijhydene.2022.05.078>.
- [183] Geim AK, Novoselov KS. The rise of graphene. *Nanoscience and Technology: A Collection of Reviews from Nature Journals* 2009:11–9. https://doi.org/10.1142/9789814287005_0002.
- [184] Bonaccorso F, Sun Z, Hasan T, Ferrari AC. Graphene photonics and optoelectronics. *Nature Photonics* 2010 4:9 2010;4:611–22. <https://doi.org/10.1038/nphoton.2010.186>.
- [185] Balandin AA, Ghosh S, Bao W, Calizo I, Teweldebrhan D, Miao F, et al. Superior thermal conductivity of single-layer graphene. *Nano Lett* 2008;8:902–7. <https://doi.org/10.1021/NL0731872>.
- [186] Lee C, Wei X, Kysar JW, Hone J. Measurement of the elastic properties and intrinsic strength of monolayer graphene. *Science (1979)* 2008;321:385–8. <https://doi.org/10.1126/SCIENCE.1157996>.
- [187] Bunch JS, Verbridge SS, Alden JS, Van Der Zande AM, Parpia JM, Craighead HG, et al. Impermeable atomic membranes from graphene sheets. *Nano Lett* 2008;8:2458–62. <https://doi.org/10.1021/NL801457B>.
- [188] Su Y, Kravets VG, Wong SL, Waters J, Geim AK, Nair RR. Impermeable barrier films and protective coatings based on reduced graphene oxide. *Nature Communications* 2014 5:1 2014;5:1–5. <https://doi.org/10.1038/ncomms5843>.
- [189] Bukola S, Beard K, Korzeniewski C, Harris JM, Creager SE. Single-Layer Graphene Sandwiched between Proton-Exchange Membranes for Selective Proton Transmission. *ACS Appl Nano Mater* 2019;2:964–74. <https://doi.org/10.1021/acsanm.8b02270>.
- [190] Graphene research, innovation and collaboration | Graphene Flagship n.d. <https://graphene-flagship.eu/> (accessed September 17, 2022).

- [191] Ferrari AC, Bonaccorso F, Fal'ko V, Novoselov KS, Roche S, Bøggild P, et al. Science and technology roadmap for graphene, related two-dimensional crystals, and hybrid systems. *Nanoscale* 2015;7:4598–810. <https://doi.org/10.1039/C4NR01600A>.
- [192] Meister M, Braun A, Hüsing B, Schmoch U, Reiss T. Graphene and other 2D materials Technology and Innovation Roadmap Version 3 2017.
- [193] Graphene Roadmap | Graphene Flagship n.d. <https://graphene-flagship.eu/innovation/industrialisation/roadmap/> (accessed September 17, 2022).
- [194] Reiss T, Hjelt K, Ferrari AC. Graphene is on track to deliver on its promises n.d. <https://doi.org/10.1038/s41565-019-0557-0>.
- [195] Liu K, Yu C, Guo W, Ni L, Yu J, Xie Y, et al. Recent research advances of self-discharge in supercapacitors: Mechanisms and suppressing strategies. *Journal of Energy Chemistry* 2021;58:94–109. <https://doi.org/10.1016/J.JECHEM.2020.09.041>.
- [196] Schmidt O, Melchior S, Hawkes A, Staffell I. Projecting the Future Levelized Cost of Electricity Storage Technologies. *Joule* 2019;3:81–100. <https://doi.org/10.1016/J.JOULE.2018.12.008>.
- [197] Prevent Blackouts With Synthetic Inertia and Supercapacitors. *Battery Tech Online* n.d. <https://www.batterytechonline.com/batteries/prevent-blackouts-with-synthetic-inertia-and-supercapacitors> (accessed June 1, 2025).
- [198] Wikipedia contributors. 2025 Iberian Peninsula blackout. *Wikipedia* n.d. https://en.wikipedia.org/wiki/2025_Iberian_Peninsula_blackout (accessed June 1, 2025).
- [199] The supercapacitor market is set to take off, and graphene could play a center role | Graphene-Info n.d. <https://www.graphene-info.com/supercapacitor-market-set-take-and-graphene-could-play-center-role> (accessed September 22, 2022).
- [200] Nybeck CN, Dodson DA, Wetz DA, Heinzl JM. Characterization of Ultracapacitors for Transient Load Applications. *IEEE Transactions on Plasma Science* 2019;47:2493–9. <https://doi.org/10.1109/TPS.2019.2904562>.
- [201] Jaan Leis, Anti Perkson, Sebastian Pohlmann, Mati Arulepp. A carbon electrode materials for double layer ultracapacitor and double layer ultracapacitor using it. EP3732699A2, 2019.
- [202] Jaan Leis, Maike Käarik, Mati Arulepp, Anti Perkson. A method for manufacturing microporous carbon particles. WO2017207593A1, 2017.
- [203] Jaan Leis, Mati Arulepp, Anti Perkson. A carbon composite electrode for the electric double-layer capacitor. EP2564404A1, 2011.
- [204] Jaan Leis, Mati Arulepp, Marko Lätt, Helle Kuura. Method of making the porous carbon material of different pore sizes and porous carbon materials produced by the method. US7803345B2, 2005.
- [205] Ultracapacitor technology | Skeleton n.d. <https://www.skeletontech.com/ultracapacitor-technology> (accessed September 24, 2022).

- [206] Skeleton Technologies Group. Company Overview: Skeleton Technologies — Global Leader in Ultracapacitor-Based Energy Storage. 2024.
- [207] Skeleton and Siemens agree on strategic partnership for the production of supercapacitors | Press | Company | Siemens n.d. <https://press.siemens.com/global/en/pressrelease/skeleton-and-siemens-agree-strategic-partnership-production-supercapacitors> (accessed September 24, 2022).
- [208] Kim T, Jung G, Yoo S, Suh KS, Ruoff RS. Activated graphene-based carbons as supercapacitor electrodes with macro- and mesopores. *ACS Nano* 2013;7:6899–905. <https://doi.org/10.1021/NN402077V>.
- [209] Li C, Zhang X, Lv Z, Wang K, Sun X, Chen X, et al. Scalable combustion synthesis of graphene-welded activated carbon for high-performance supercapacitors. *Chemical Engineering Journal* 2021;414:128781. <https://doi.org/10.1016/J.CEJ.2021.128781>.
- [210] Stoller MD, Park S, Yanwu Z, An J, Ruoff RS. Graphene-Based ultracapacitors. *Nano Lett* 2008;8:3498–502. <https://doi.org/10.1021/NL802558Y>.
- [211] Pohlmann S. Metrics and methods for moving from research to innovation in energy storage. *Nature Communications* 2022 13:1 2022;13:1–5. <https://doi.org/10.1038/s41467-022-29257-w>.
- [212] Li Z, Gadipelli S, Li H, Howard CA, Brett DJL, Shearing PR, et al. Tuning the interlayer spacing of graphene laminate films for efficient pore utilization towards compact capacitive energy storage. *Nature Energy* 2020 5:2 2020;5:160–8. <https://doi.org/10.1038/s41560-020-0560-6>.
- [213] Davies A, Yu A. Material advancements in supercapacitors: From activated carbon to carbon nanotube and graphene. *Can J Chem Eng* 2011;89:1342–57. <https://doi.org/10.1002/CJCE.20586>.
- [214] Li Z, Gadipelli S, Yang Y, He G, Guo J, Li J, et al. Exceptional supercapacitor performance from optimized oxidation of graphene-oxide. *Energy Storage Mater* 2019;17:12–21. <https://doi.org/10.1016/J.ENSM.2018.12.006>.
- [215] Dang F, Yang P, Zhao W, Liu JZ, Wu H, Liu A, et al. Tuning capacitance of graphene films via a robust routine of adjusting their hierarchical structures. *Electrochim Acta* 2019;298:254–64. <https://doi.org/10.1016/J.ELECTACTA.2018.12.071>.
- [216] Nithya VD. A review on holey graphene electrode for supercapacitor. *J Energy Storage* 2021;44:103380. <https://doi.org/10.1016/J.EST.2021.103380>.
- [217] Chen Y, Zhang X, Zhang D, Yu P, Ma Y. High performance supercapacitors based on reduced graphene oxide in aqueous and ionic liquid electrolytes. *Carbon N Y* 2011;49:573–80. <https://doi.org/10.1016/J.CARBON.2010.09.060>.
- [218] Dong X, Hu N, Wei L, Su Y, Wei H, Yao L, et al. A new strategy to prepare N-doped holey graphene for high-volumetric supercapacitors. *J Mater Chem A Mater* 2016;4:9739–43. <https://doi.org/10.1039/C6TA01406B>.

- [219] Zhu G, Xi C, Liu Y, Zhu J, Shen X. CN foam loaded with few-layer graphene nanosheets for high-performance supercapacitor electrodes. *J Mater Chem A Mater* 2015;3:7591–9. <https://doi.org/10.1039/C5TA00837A>.
- [220] Choi BG, Yang M, Hong WH, Choi JW, Huh YS. 3D macroporous graphene frameworks for supercapacitors with high energy and power densities. *ACS Nano* 2012;6:4020–8. <https://doi.org/10.1021/NN3003345>.
- [221] Lee JS, Kim SI, Yoon JC, Jang JH. Chemical vapor deposition of mesoporous graphene nanoballs for supercapacitor. *ACS Nano* 2013;7:6047–55. https://doi.org/10.1021/NN401850Z/SUPPL_FILE/NN401850Z_SI_001.PDF.
- [222] Lei Z, Lu L, Zhao XS. The electrocapacitive properties of graphene oxide reduced by urea. *Energy Environ Sci* 2012;5:6391–9. <https://doi.org/10.1039/C1EE02478G>.
- [223] Wu ZS, Winter A, Chen L, Sun Y, Turchanin A, Feng X, et al. Three-Dimensional Nitrogen and Boron Co-doped Graphene for High-Performance All-Solid-State Supercapacitors. *Advanced Materials* 2012;24:5130–5. <https://doi.org/10.1002/ADMA.201201948>.
- [224] Liu Z, Jiang L, Sheng L, Zhou Q, Wei T, Zhang B, et al. Oxygen Clusters Distributed in Graphene with “Paddy Land” Structure: Ultrahigh Capacitance and Rate Performance for Supercapacitors. *Adv Funct Mater* 2018;28:1705258. <https://doi.org/10.1002/ADFM.201705258>.
- [225] Sui ZY, Meng YN, Xiao PW, Zhao ZQ, Wei ZX, Han BH. Nitrogen-doped graphene aerogels as efficient supercapacitor electrodes and gas adsorbents. *ACS Appl Mater Interfaces* 2015;7:1431–8. <https://doi.org/10.1021/AM5042065>.
- [226] Haque E, Islam MM, Pourazadi E, Hassan M, Faisal SN, Roy AK, et al. Nitrogen doped graphene via thermal treatment of composite solid precursors as a high performance supercapacitor. *RSC Adv* 2015;5:30679–86. <https://doi.org/10.1039/C4RA17262K>.
- [227] Lu Y, Zhang F, Zhang T, Leng K, Zhang L, Yang X, et al. Synthesis and supercapacitor performance studies of N-doped graphene materials using o-phenylenediamine as the double-N precursor. *Carbon N Y* 2013;63:508–16. <https://doi.org/10.1016/J.CARBON.2013.07.026>.
- [228] Shao Y, El-Kady MF, Lin CW, Zhu G, Marsh KL, Hwang JY, et al. 3D Freeze-Casting of Cellular Graphene Films for Ultrahigh-Power-Density Supercapacitors. *Adv Mater* 2016;28:6719–26. <https://doi.org/10.1002/ADMA.201506157>.
- [229] Jhajharia SK, Selvaraj K. Non-templated ambient nanoporation of graphene: a novel scalable process and its exploitation for energy and environmental applications. *Nanoscale* 2015;7:19705–13. <https://doi.org/10.1039/C5NR05715A>.
- [230] Peng YY, Liu YM, Chang JK, Wu CH, Ger M Der, Pu NW, et al. A facile approach to produce holey graphene and its application in supercapacitors. *Carbon N Y* 2015;81:347–56. <https://doi.org/10.1016/J.CARBON.2014.09.067>.
- [231] Jiang ZJ, Jiang Z, Chen W. The role of holes in improving the performance of nitrogen-doped holey graphene as an active electrode material for supercapacitor and oxygen reduction reaction. *J Power Sources* 2014;251:55–65. <https://doi.org/10.1016/J.JPOWSOUR.2013.11.031>.

- [232] Chen Y, Li Y, Yao F, Peng C, Cao C, Feng Y, et al. Nitrogen and fluorine co-doped holey graphene hydrogel as a binder-free electrode material for flexible solid-state supercapacitors. *Sustain Energy Fuels* 2019;3:2237–45. <https://doi.org/10.1039/C9SE00142E>.
- [233] Nazarian-Samani M, Haghghat-Shishavan S, Nazarian-Samani M, Kim MS, Cho BW, Oh SH, et al. Rational hybrid modulation of P, N dual-doped holey graphene for high-performance supercapacitors. *J Power Sources* 2017;372:286–96. <https://doi.org/10.1016/J.JPOWSOUR.2017.10.087>.
- [234] Jia S, Zang J, Tian P, Zhou S, Cai H, Tian X, et al. A 3-D covalently crosslinked N-doped porous carbon/holey graphene composite for quasi-solid-state supercapacitors. *Microporous and Mesoporous Materials* 2020;293:109796. <https://doi.org/10.1016/J.MICROMESO.2019.109796>.
- [235] Fan Z, Zhu J, Sun X, Cheng Z, Liu Y, Wang Y. High Density of Free-Standing Holey Graphene/PPy Films for Superior Volumetric Capacitance of Supercapacitors. *ACS Appl Mater Interfaces* 2017;9:21763–72. <https://doi.org/10.1021/ACSAMI.7B03477>.
- [236] Chen N, Ni L, Zhou J, Zhu G, Kang Q, Zhang Y, et al. Sandwich-Like Holey Graphene/PANI/Graphene Nanohybrid for Ultrahigh-Rate Supercapacitor. *ACS Appl Energy Mater* 2018;1:5189–97. <https://doi.org/10.1021/ACSAEM.8B00725>.
- [237] Xu Y, Sheng K, Li C, Shi G. Self-assembled graphene hydrogel via a one-step hydrothermal process. *ACS Nano* 2010;4:4324–30. <https://doi.org/10.1021/NN101187Z>.
- [238] Wu ZS, Sun Y, Tan YZ, Yang S, Feng X, Müllen K. Three-dimensional graphene-based macro- and mesoporous frameworks for high-performance electrochemical capacitive energy storage. *J Am Chem Soc* 2012;134:19532–5. <https://doi.org/10.1021/JA308676H>.
- [239] Xu Y, Lin Z, Zhong X, Huang X, Weiss NO, Huang Y, et al. Holey graphene frameworks for highly efficient capacitive energy storage. *Nature Communications* 2014 5:1 2014;5:1–8. <https://doi.org/10.1038/ncomms5554>.
- [240] Xu Y, Chen CY, Zhao Z, Lin Z, Lee C, Xu X, et al. Solution Processable Holey Graphene Oxide and Its Derived Macrostructures for High-Performance Supercapacitors. *Nano Lett* 2015;15:4605–10. <https://doi.org/10.1021/ACS.NANOLETT.5B01212>.
- [241] Lv W, Tang DM, He YB, You CH, Shi ZQ, Chen XC, et al. Low-temperature exfoliated graphenes: Vacuum-promoted exfoliation and electrochemical energy storage. *ACS Nano* 2009;3:3730–6. <https://doi.org/10.1021/NN900933U>.
- [242] Yoon Y, Lee K, Baik C, Yoo H, Min M, Park Y, et al. Anti-Solvent Derived Non-Stacked Reduced Graphene Oxide for High Performance Supercapacitors. *Advanced Materials* 2013;25:4437–44. <https://doi.org/10.1002/ADMA.201301230>.
- [243] Tao Y, Xie X, Lv W, Tang DM, Kong D, Huang Z, et al. Towards ultrahigh volumetric capacitance: graphene derived highly dense but porous carbons for supercapacitors. *Scientific Reports* 2013 3:1 2013;3:1–8. <https://doi.org/10.1038/srep02975>.

- [244] Lota G, Centeno TA, Frackowiak E, Stoeckli F. Improvement of the structural and chemical properties of a commercial activated carbon for its application in electrochemical capacitors. *Electrochim Acta* 2008;53:2210–6. <https://doi.org/10.1016/J.ELECTACTA.2007.09.028>.
- [245] Bonaccorso F, Colombo L, Yu G, Stoller M, Tozzini V, Ferrari AC, et al. Graphene, related two-dimensional crystals, and hybrid systems for energy conversion and storage. *Science* (1979) 2015;347. <https://doi.org/10.1126/SCIENCE.1246501>.
- [246] Joy R, Balakrishnan NTM, Das A, Shafeek S, Thakur VK, Zaghib K, et al. Graphene: Chemistry and Applications for Lithium-Ion Batteries. *Electrochem* 2022, Vol 3, Pages 143-183 2022;3:143–83. <https://doi.org/10.3390/ELECTROCHEM3010010>.
- [247] Pan Z, Brett DJL, He G, Parkin IP, Pan Z, He G, et al. Progress and Perspectives of Organosulfur for Lithium–Sulfur Batteries. *Adv Energy Mater* 2022;12:2103483. <https://doi.org/10.1002/AENM.202103483>.
- [248] Zhang W, Wu Y, Xu Z, Li H, Xu M, Li J, et al. Rationally Designed Sodium Chromium Vanadium Phosphate Cathodes with Multi-Electron Reaction for Fast-Charging Sodium-Ion Batteries. *Adv Energy Mater* 2022;12:2201065. <https://doi.org/10.1002/AENM.202201065>.
- [249] Han P, Yue Y, Liu Z, Xu W, Zhang L, Xu H, et al. Graphene oxide nanosheets/multi-walled carbon nanotubes hybrid as an excellent electrocatalytic material towards VO_2^+/VO_2^+ redox couples for vanadium redox flow batteries. *Energy Environ Sci* 2011;4:4710–7. <https://doi.org/10.1039/C1EE01776D>.
- [250] Arboretti R, Ceccato R, Pegoraro L, Salmaso L. Design of Experiments and machine learning for product innovation: A systematic literature review. *Qual Reliab Eng Int* 2022;38:1131–56. <https://doi.org/10.1002/QRE.3025>.
- [251] Fisher RA. Design of Experiments. *Br Med J* 1936;1:554.
- [252] Box GEP, Wilson KB. On the Experimental Attainment of Optimum Conditions. *Journal of the Royal Statistical Society: Series B (Methodological)* 1951;13:1–38. <https://doi.org/10.1111/J.2517-6161.1951.TB00067.X>.
- [253] Taguchi G. System of experimental design; engineering methods to optimize quality and minimize costs. 1987.
- [254] McKay MD, Beckman RJ, Conover WJ. A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code. *Technometrics* 1979;21:239. <https://doi.org/10.2307/1268522>.
- [255] Peng J, Schwalbe-Koda D, Akkiraju K, Xie T, Giordano L, Yu Y, et al. Human- and machine-centred designs of molecules and materials for sustainability and decarbonization. *Nature Reviews Materials* 2022 2022:1–19. <https://doi.org/10.1038/s41578-022-00466-5>.
- [256] Seko A, Togo A, Hayashi H, Tsuda K, Chaput L, Tanaka I. Prediction of Low-Thermal-Conductivity Compounds with First-Principles Anharmonic Lattice-Dynamics Calculations and Bayesian Optimization. *Phys Rev Lett* 2015;115:205901. <https://doi.org/10.1103/PHYSREVLETT.115.205901>.

- [257] Home | ICSD n.d. <https://icsd.products.fiz-karlsruhe.de/> (accessed August 28, 2022).
- [258] Inorganic Material Database (AtomWork) - DICE :: National Institute for Materials Science n.d. <https://crystdb.nims.go.jp/en/> (accessed August 28, 2022).
- [259] Materials Project - Home n.d. <https://materialsproject.org/> (accessed August 28, 2022).
- [260] Sun S, Tiihonen A, Oviedo F, Liu Z, Thapa J, Zhao Y, et al. A data fusion approach to optimize compositional stability of halide perovskites. *Matter* 2021;4:1305–22. <https://doi.org/10.1016/J.MATT.2021.01.008>.
- [261] Zheng B, Zheng Z, Gu GX. Uncertainty quantification and prediction for mechanical properties of graphene aerogels via Gaussian process metamodels. *Nano Futures* 2021;5. <https://doi.org/10.1088/2399-1984/ac3c8f>.
- [262] Schmitz G, Schnieder B. Adaptive regularized Gaussian process regression for application in the context of hydrogen adsorption on graphene sheets. *J Comput Chem* 2023;44:732–44. <https://doi.org/10.1002/jcc.27035>.
- [263] Yamawaki M, Ohnishi M, Ju S, Shiomi J. Multifunctional structural design of graphene thermoelectrics by Bayesian optimization. *Sci Adv* 2018;4. <https://doi.org/10.1126/SCIADV.AAR4192>.
- [264] GitHub - tsudalab/combo: COMmon Bayesian Optimization n.d. <https://github.com/tsudalab/combo> (accessed August 27, 2022).
- [265] Chaney LE, van Beek A, Downing JR, Zhang J, Zhang H, Hui J, et al. Bayesian Optimization of Environmentally Sustainable Graphene Inks Produced by Wet Jet Milling. *Small* 2024;20. <https://doi.org/10.1002/sml.202309579>.
- [266] Patel AG, Johnson L, Arroyave R, Lutkenhaus JL. Design of multifunctional supercapacitor electrodes using an informatics approach. *Mol Syst Des Eng* 2019;4:654–63. <https://doi.org/10.1039/C8ME00060C>.
- [267] scikit-learn: machine learning in Python — scikit-learn 1.1.3 documentation n.d. <https://scikit-learn.org/stable/> (accessed December 3, 2022).
- [268] Alghfeli A, Fisher TS. Sequential Bayesian-optimized graphene synthesis by direct solar-thermal chemical vapor deposition. *Sci Rep* 2024;14. <https://doi.org/10.1038/s41598-024-54005-z>.
- [269] Wahab H, Jain V, Tyrrell AS, Seas MA, Kotthoff L, Johnson PA. Machine-learning-assisted fabrication: Bayesian optimization of laser-induced graphene patterning using in-situ Raman analysis. *Carbon N Y* 2020;167:609–19. <https://doi.org/10.1016/J.CARBON.2020.05.087>.
- [270] GitHub - mlr-org/mlrMBO: Toolbox for Bayesian Optimization and Model-Based Optimization in R n.d. <https://github.com/mlr-org/mlrMBO> (accessed August 27, 2022).
- [271] Ebikade EO, Wang Y, Samulewicz N, Hasa B, Vlachos D. Active learning-driven quantitative synthesis–structure–property relations for improving performance and revealing active sites of nitrogen-doped carbon for the hydrogen evolution reaction. *React Chem Eng* 2020;5:2134–47. <https://doi.org/10.1039/D0RE00243G>.

- [272] GitHub - capaulson/pyKriging: Welcome to the User Friendly Python Kriging Toolbox! n.d. <https://github.com/capaulson/pyKriging> (accessed August 27, 2022).
- [273] Cross EJ, Rogers TJ, Pitchforth DJ, Gibson SJ, Zhang S, Jones MR. A spectrum of physics-informed Gaussian processes for regression in engineering. *Data-Centric Engineering* 2024;5. <https://doi.org/10.1017/dce.2024.2>.
- [274] Raissi M, Perdikaris P, Ahmadi N, Karniadakis GE. *Physics-Informed Neural Networks and Extensions*. 2024.
- [275] Pan R, Gu M, Wu J. Physics-informed Gaussian process regression of in operando capacitance for carbon supercapacitors. *Energy Advances* 2023;2:843–53. <https://doi.org/10.1039/D3YA00071K>.
- [276] Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. *Nature* 2018 559:7715 2018;559:547–55. <https://doi.org/10.1038/s41586-018-0337-2>.
- [277] Atha DJ, Jahanshahi MR. Evaluation of deep learning approaches based on convolutional neural networks for corrosion detection. *Struct Health Monit* 2017;17:1110–28. <https://doi.org/10.1177/1475921717737051>.
- [278] Oh E, Liu R, Nel A, Gemill KB, Bilal M, Cohen Y, et al. Meta-analysis of cellular toxicity for cadmium-containing quantum dots. *Nature Nanotechnology* 2016 11:5 2016;11:479–86. <https://doi.org/10.1038/nnano.2015.338>.
- [279] Najjar IMR, Sadoun AM, Alsuruji GS, Elaziz MA, Wagih A. Predicting the mechanical properties of Cu–Al₂O₃ nanocomposites using machine learning and finite element simulation of indentation experiments. *Ceram Int* 2022;48:7748–58. <https://doi.org/10.1016/J.CERAMINT.2021.11.322>.
- [280] Hruska E, Gale A, Liu F. Bridging the Experiment-Calculation Divide: Machine Learning Corrections to Redox Potential Calculations in Implicit and Explicit Solvent Models. *J Chem Theory Comput* 2022;18:1096–108. <https://doi.org/10.1021/ACS.JCTC.1C01040>.
- [281] Schattauer C, Todorović M, Ghosh K, Rinke P, Libisch F. Machine learning sparse tight-binding parameters for defects. *Npj Computational Materials* 2022 8:1 2022;8:1–11. <https://doi.org/10.1038/s41524-022-00791-x>.
- [282] Zhu S, Li J, Ma L, He C, Liu E, He F, et al. Artificial neural network enabled capacitance prediction for carbon-based supercapacitors. *Mater Lett* 2018;233:294–7. <https://doi.org/10.1016/j.matlet.2018.09.028>.
- [283] Su H, Lin S, Deng S, Lian C, Shang Y, Liu H. Predicting the capacitance of carbon-based electric double layer capacitors by machine learning. *Nanoscale Adv* 2019;1:2162–6. <https://doi.org/10.1039/c9na00105k>.
- [284] Zhou M, Gallegos A, Liu K, Dai S, Wu J. Insights from machine learning of carbon electrodes for electric double layer capacitors. *Carbon N Y* 2020;157:147–52. <https://doi.org/10.1016/j.carbon.2019.08.090>.
- [285] Zhou M, Vassallo A, Wu J. Data-Driven Approach to Understanding the In-Operando Performance of Heteroatom-Doped Carbon Electrodes. *ACS Appl Energy Mater* 2020;3:5993–6000. <https://doi.org/10.1021/acsaem.0c01059>.

- [286] Gheyntanzadeh M, Baghban A, Habibzadeh S, Mohaddespour A, Abida O. Insights into the estimation of capacitance for carbon-based supercapacitors. *RSC Adv* 2021;11:5479–86. <https://doi.org/10.1039/D0RA09837J>.
- [287] Liu P, Wen Y, Huang L, Zhu X, Wu R, Ai S, et al. An emerging machine learning strategy for the assisted-design of high-performance supercapacitor materials by mining the relationship between capacitance and structural features of porous carbon. *Journal of Electroanalytical Chemistry* 2021;899. <https://doi.org/10.1016/j.jelechem.2021.115684>.
- [288] Chen T, Guestrin C. Xgboost: A scalable tree boosting system. *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 2016, p. 785–94.
- [289] Mishra S, Srivastava R, Muhammad A, Amit A, Chiavazzo E, Fasano M, et al. The impact of physicochemical features of carbon electrodes on the capacitive performance of supercapacitors: a machine learning approach. *Sci Rep* 2023;13. <https://doi.org/10.1038/s41598-023-33524-1>.
- [290] Wang T, Pan R, Martins ML, Cui J, Huang Z, Thapaliya BP, et al. Machine-learning-assisted material discovery of oxygen-rich highly porous carbon active materials for aqueous supercapacitors. *Nat Commun* 2023;14. <https://doi.org/10.1038/s41467-023-40282-1>.
- [291] Zhang L, Zhang F, Yang X, Long G, Wu Y, Zhang T, et al. Porous 3D graphene-based bulk materials with exceptional high surface area and excellent conductivity for supercapacitors. *Sci Rep* 2013;3. <https://doi.org/10.1038/srep01408>.
- [292] Chen S, Wang W, Zhang X, Wang X. High-Performance Supercapacitors Based on Graphene/Activated Carbon Hybrid Electrodes Prepared via Dry Processing. *Batteries* 2024;10. <https://doi.org/10.3390/batteries10060195>.
- [293] Mousavi SM, Hashemi SA, Kalashgrani MY, Gholami A, Binazadeh M, Chiang WH, et al. Recent advances in energy storage with graphene oxide for supercapacitor technology. *Sustain Energy Fuels* 2023;7:5176–97. <https://doi.org/10.1039/d3se00867c>.
- [294] İkinci M, Zehni OC, Bissett MA. Structural Supercapacitors Based on Graphene Nanoplatelet-Coated Carbon Fiber Electrodes. *ACS Appl Nano Mater* 2024. <https://doi.org/10.1021/acsanm.4c04801>.
- [295] Saad AG, Emad-Eldeen A, Tawfik WZ, El-Deen AG. Data-driven machine learning approach for predicting the capacitance of graphene-based supercapacitor electrodes. *J Energy Storage* 2022;55. <https://doi.org/10.1016/j.est.2022.105411>.
- [296] Wang J, Yang J, Huang T, Yin W. Mn_{0.5}Co_{2.5}O₄ nanofibers sandwiched in graphene sheets for efficient supercapacitor electrode materials. *RSC Adv* 2016;6:103923–9. <https://doi.org/10.1039/c6ra21281f>.
- [297] Chang H, Kang J, Chen L, Wang J, Ohmura K, Chen N, et al. Low-temperature solution-processable Ni(OH)₂ ultrathin nanosheet/N-graphene nanohybrids for high-performance supercapacitor electrodes. *Nanoscale* 2014;6:5960–6. <https://doi.org/10.1039/c4nr00655k>.
- [298] Chenwittayakhachon A, Jitapunkul K, Nakpalad B, Worrayotkovit P, Namuangruk S, Sirisinudomkit P, et al. Machine learning approach to understanding the

- 'synergistic' pseudocapacitive effects of heteroatom doped graphene. *2d Mater* 2023;10. <https://doi.org/10.1088/2053-1583/acaf8d>.
- [299] Helmholtz H. Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern mit Anwendung auf die thierisch-elektrischen Versuche. *Ann Phys* 1853;165:211–33. <https://doi.org/10.1002/andp.18531650603>.
- [300] Gouy M. Sur la constitution de la charge électrique à la surface d'un électrolyte. *J Phys Theor Appl* 1910;9:457–68.
- [301] Chapman DL. LI. A contribution to the theory of electrocapillarity. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 1913;25:475–81.
- [302] Stern O. Zur theorie der elektrolytischen doppelschicht. *Zeitschrift Für Elektrochemie Und Angewandte Physikalische Chemie* 1924;30:508–16.
- [303] Deshsorn K, Payakkachon K, Chairsithong T, Jitapunkul K, Lawtrakul L, lamprasertkun P. Unlocking the Full Potential of Heteroatom-Doped Graphene-Based Supercapacitors through Stacking Models and SHAP-Guided Optimization. *J Chem Inf Model* 2023;63:5077–88. <https://doi.org/10.1021/acs.jcim.3c00670>.
- [304] Deshsorn K, Sirisaksoontorn W, Hirunpinyopas W, lamprasertkun P. Exploring how base model combination affects the results of a “stacking” ensemble machine learning model: An applied study on optimization of heteroatom doped carbon data. *FlatChem* 2025;50. <https://doi.org/10.1016/j.flatc.2025.100827>.
- [305] Kaya D, Koroglu D, Aydın E, Uralcan B. Optimization of capacitance in supercapacitors by constructing an experimentally validated hybrid artificial neural networks-genetic algorithm framework. *J Power Sources* 2023;568. <https://doi.org/10.1016/j.jpowsour.2023.232987>.
- [306] Novoselov KS, Fal'Ko VI, Colombo L, Gellert PR, Schwab MG, Kim K. A roadmap for graphene. *Nature* 2012 490:7419 2012;490:192–200. <https://doi.org/10.1038/nature11458>.
- [307] Galande C, Gao W, Mathkar A, Dattelbaum AM, Narayanan TN, Mohite AD, et al. Science and engineering of graphene oxide. *Particle and Particle Systems Characterization* 2014;31:619–38. <https://doi.org/10.1002/PPSC.201300232>.
- [308] Joshi DJ, Koduru JR, Malek NI, Hussain CM, Kailasa SK. Surface modifications and analytical applications of graphene oxide: A review. *TrAC Trends in Analytical Chemistry* 2021;144:116448. <https://doi.org/10.1016/J.TRAC.2021.116448>.
- [309] Li D, Müller MB, Gilje S, Kaner RB, Wallace GG. Processable aqueous dispersions of graphene nanosheets. *Nature Nanotechnology* 2008 3:2 2008;3:101–5. <https://doi.org/10.1038/nnano.2007.451>.
- [310] About Us - The Sixth Element (Changzhou) Materials Technology Co.,Ltd n.d. <https://www.c6th.com/about-us> (accessed October 7, 2022).
- [311] Abraham J, Vasu KS, Williams CD, Gopinadhan K, Su Y, Cherian CT, et al. Tunable sieving of ions using graphene oxide membranes. *Nature Nanotechnology* 2017 12:6 2017;12:546–50. <https://doi.org/10.1038/nnano.2017.21>.

- [312] Bhaskar S, Visweswar Kambhampati NS, Ganesh KM, Pa MS, Srinivasan V, Ramamurthy SS. Metal-Free, Graphene Oxide-Based Tunable Soliton and Plasmon Engineering for Biosensing Applications. *ACS Appl Mater Interfaces* 2021;13:17046–61. <https://doi.org/10.1021/ACSAMI.1C01024>.
- [313] Kumar SSA, Bashir S, Ramesh K, Ramesh S. New perspectives on Graphene/Graphene oxide based polymer nanocomposites for corrosion applications: The relevance of the Graphene/Polymer barrier coatings. *Prog Org Coat* 2021;154:106215. <https://doi.org/10.1016/J.PORGCOAT.2021.106215>.
- [314] El-Kady MF, Strong V, Dubin S, Kaner RB. Laser scribing of high-performance and flexible graphene-based electrochemical capacitors. *Science* (1979) 2012;335:1326–30. <https://doi.org/10.1126/SCIENCE.1216744>.
- [315] Akhavan O, Ghaderi E. Escherichia coli bacteria reduce graphene oxide to bactericidal graphene in a self-limiting manner. *Carbon N Y* 2012;50:1853–60. <https://doi.org/10.1016/J.CARBON.2011.12.035>.
- [316] Gadipelli S, Lu Y, Skipper NT, Yildirim T, Guo Z. Design of hyperporous graphene networks and their application in solid-amine based carbon capture systems. *J Mater Chem A Mater* 2017;5:17833–40. <https://doi.org/10.1039/c7ta05789j>.
- [317] Zhu Y, Murali S, Stoller MD, Velamakanni A, Piner RD, Ruoff RS. Microwave assisted exfoliation and reduction of graphite oxide for ultracapacitors. *Carbon N Y* 2010;48:2118–22. <https://doi.org/10.1016/J.CARBON.2010.02.001>.
- [318] Chen W, Yan L, Bangal PR. Preparation of graphene by the rapid and mild thermal reduction of graphene oxide induced by microwaves. *Carbon N Y* 2010;48:1146–52. <https://doi.org/10.1016/J.CARBON.2009.11.037>.
- [319] Guex LG, Sacchi B, Peuvot KF, Andersson RL, Pourrahimi AM, Ström V, et al. Experimental review: chemical reduction of graphene oxide (GO) to reduced graphene oxide (rGO) by aqueous chemistry. *Nanoscale* 2017;9:9562–71. <https://doi.org/10.1039/C7NR02943H>.
- [320] Zhou Y, Bao Q, Tang LAL, Zhong Y, Loh KP. Hydrothermal dehydration for the “green” reduction of exfoliated graphene oxide to graphene and demonstration of tunable optical limiting properties. *Chemistry of Materials* 2009;21:2950–6. <https://doi.org/10.1021/CM9006603>.
- [321] Toh SY, Loh KS, Kamarudin SK, Daud WRW. Graphene production via electrochemical reduction of graphene oxide: Synthesis and characterisation. *Chemical Engineering Journal* 2014;251:422–34. <https://doi.org/10.1016/J.CEJ.2014.04.004>.
- [322] Stankovich S, Dikin DA, Piner RD, Kohlhaas KA, Kleinhammes A, Jia Y, et al. Synthesis of graphene-based nanosheets via chemical reduction of exfoliated graphite oxide. *Carbon N Y* 2007;45:1558–65. <https://doi.org/10.1016/J.CARBON.2007.02.034>.
- [323] Shin HJ, Kim KK, Benayad A, Yoon SM, Park HK, Jung IS, et al. Efficient Reduction of Graphite Oxide by Sodium Borohydride and Its Effect on Electrical Conductance. *Adv Funct Mater* 2009;19:1987–92. <https://doi.org/10.1002/ADFM.200900167>.

- [324] Zhang J, Yang H, Shen G, Cheng P, Zhang J, Guo S. Reduction of graphene oxide via ascorbic acid. *Chemical Communications* 2010;46:1112–4. <https://doi.org/10.1039/B917705A>.
- [325] Ismail Z. Green reduction of graphene oxide by plant extracts: A short review. *Ceram Int* 2019;45:23857–68. <https://doi.org/10.1016/J.CERAMINT.2019.08.114>.
- [326] Vu TT, Suarez Sanchez R, Noriega Perez D. A method for the manufacture of graphene oxide from expanded kish graphite. 2020229881A1, 2019.
- [327] Vu TT, Suarez Sanchez R. A method for the manufacture of reduced graphene oxide from expanded kish graphite. 2020229882A1, 2019.
- [328] Vu TT, Perez Vidal O, Suarez Sanchez R. A method for the manufacture of pristine graphene from kish graphite. 2019220174A1, 2018.
- [329] Cyclic Voltammetry - Chemistry LibreTexts n.d. [https://chem.libretexts.org/Bookshelves/Analytical_Chemistry/Supplemental_Modules_\(Analytical_Chemistry\)/Instrumental_Analysis/Cyclic_Voltammetry](https://chem.libretexts.org/Bookshelves/Analytical_Chemistry/Supplemental_Modules_(Analytical_Chemistry)/Instrumental_Analysis/Cyclic_Voltammetry) (accessed October 9, 2022).
- [330] Elgrishi N, Rountree KJ, McCarthy BD, Rountree ES, Eisenhart TT, Dempsey JL. A Practical Beginner's Guide to Cyclic Voltammetry. *J Chem Educ* 2018;95:197–206. <https://doi.org/10.1021/ACS.JCHEMED.7B00361>.
- [331] Pajkossy T, Kolb DM. Anion-adsorption-related frequency-dependent double layer capacitance of the platinum-group metals in the double layer region. *Electrochim Acta* 2008;53:7403–9. <https://doi.org/10.1016/J.ELECTACTA.2007.11.068>.
- [332] Stoller MD, Ruoff RS. Best practice methods for determining an electrode material's performance for ultracapacitors. *Energy Environ Sci* 2010;3:1294–301. <https://doi.org/10.1039/C0EE00074D>.
- [333] Nikolic J, Expósito E, Iniesta J, González-García J, Montiel V. Theoretical Concepts and Applications of a Rotating Disk Electrode. *J Chem Educ* 2000;77:1191–4. <https://doi.org/10.1021/ED077P1191>.
- [334] Zhang S, Pan N. Supercapacitors Performance Evaluation. *Adv Energy Mater* 2015;5:1401401. <https://doi.org/10.1002/AENM.201401401>.
- [335] Shao Y, El-Kady MF, Sun J, Li Y, Zhang Q, Zhu M, et al. Design and Mechanisms of Asymmetric Supercapacitors. *Chem Rev* 2018;118:9233–80. <https://doi.org/10.1021/ACS.CHEMREV.8B00252>.
- [336] Ferrari AC, Basko DM. Raman spectroscopy as a versatile tool for studying the properties of graphene. *Nature Nanotechnology* 2013 8:4 2013;8:235–46. <https://doi.org/10.1038/nnano.2013.46>.
- [337] Ferrari AC. Raman spectroscopy of graphene and graphite: Disorder, electron-phonon coupling, doping and nonadiabatic effects. *Solid State Commun* 2007;143:47–57. <https://doi.org/10.1016/J.SSC.2007.03.052>.
- [338] Stankovich S, Dikin DA, Piner RD, Kohlhaas KA, Kleinhammes A, Jia Y, et al. Synthesis of graphene-based nanosheets via chemical reduction of exfoliated graphite oxide. *Carbon N Y* 2007;45:1558–65. <https://doi.org/10.1016/J.CARBON.2007.02.034>.

- [339] Kauling AP, Seefeldt AT, Pisoni DP, Pradeep RC, Bentini R, Oliveira RVB, et al. The Worldwide Graphene Flake Production. *Advanced Materials* 2018;30. <https://doi.org/10.1002/ADMA.201803784>.
- [340] Characterisation of the structure of graphene GPG145 - NPL n.d. <https://www.npl.co.uk/gpgs/graphene-structure-characterisation> (accessed October 12, 2022).
- [341] Marchesini S, Turner P, Paton KR, Reed BP, Brennan B, Koziol K, et al. Gas physisorption measurements as a quality control tool for the properties of graphene/graphite powders. *Carbon* N Y 2020;167:585–95. <https://doi.org/10.1016/J.CARBON.2020.05.083>.
- [342] Murphy KP. *Machine Learning: A Probabilistic Perspective*. The MIT Press; 2012.
- [343] Pinheiro M, Ge F, Ferré N, Dral PO, Barbatti M. Choosing the right molecular machine learning potential. *Chem Sci* 2021;12:14396–413. <https://doi.org/10.1039/D1SC03564A>.
- [344] What is the Difference Between a Parameter and a Hyperparameter? n.d. <https://machinelearningmastery.com/difference-between-a-parameter-and-a-hyperparameter/> (accessed October 22, 2022).
- [345] 1.10. Decision Trees — scikit-learn 1.1.2 documentation n.d. <https://scikit-learn.org/stable/modules/tree.html> (accessed October 22, 2022).
- [346] 1.11. Ensembles: Gradient boosting, random forests, bagging, voting, stacking — scikit-learn 1.8.0 documentation n.d. <https://scikit-learn.org/stable/modules/ensemble.html#random-forests> (accessed December 30, 2025).
- [347] Garnett R. *Bayesian optimization*. Cambridge University Press; 2023.
- [348] Williams CKI, Rasmussen CE. *Gaussian processes for machine learning*. vol. 2. MIT press Cambridge, MA; 2006.
- [349] Gaussian Process Regressor — scikit-learn 1.1.3 documentation n.d. https://scikit-learn.org/stable/modules/generated/sklearn.gaussian_process.GaussianProcessRegressor.html#sklearn.gaussian_process.GaussianProcessRegressor (accessed December 3, 2022).
- [350] 1.7. Gaussian Processes — scikit-learn 1.1.3 documentation n.d. https://scikit-learn.org/stable/modules/gaussian_process.html (accessed December 3, 2022).
- [351] *Bayesian Optimization Book* | Copyright 2021 Roman Garnett, to be published by Cambridge University Press n.d. <https://bayesoptbook.com/> (accessed July 17, 2022).
- [352] Coleman HW, Steele WG. *Experimentation, validation, and uncertainty analysis for engineers: Fourth edition*. 2018. <https://doi.org/10.1002/9781119417989>.
- [353] Garsany Y, Singer IL, Swider-Lyons KE. Impact of film drying procedures on RDE characterization of Pt/VC electrocatalysts. *Journal of Electroanalytical Chemistry* 2011;662:396–406. <https://doi.org/10.1016/j.jelechem.2011.09.016>.

- [354] Garsany Y, Baturina OA, Swider-Lyons KE, Kocha SS. Experimental methods for quantifying the activity of platinum electrocatalysts for the oxygen reduction reaction. *Anal Chem* 2010;82:6321–8. <https://doi.org/10.1021/ac100306c>.
- [355] Okubo K, Thik J, Yamaguchi T, Ling C. Computer vision enabled high-quality electrochemical experimentation. *Digital Discovery* 2024;3:2183–91. <https://doi.org/10.1039/d4dd00213j>.
- [356] Chen W, Cui H-W, Liao L-W, Xu Y-J, Cai J, Chen Y-X. Challenges in Unravelling the Intrinsic Kinetics of Gas Reactions at Rotating Disk Electrodes by Koutecky-Levich Equation. *Journal of Physical Chemistry C* 2023;127:16235–48. <https://doi.org/10.1021/acs.jpcc.3c02677>.
- [357] Bhandari S, Narangoda P V., Mogensen SO, Tesch MF, Mechler AK. Effect of Experimental Parameters on the Electrocatalytic Performance in Rotating Disc Electrode Measurements: Case Study of Oxygen Evolution on Ni-Co-Oxide in Alkaline Media. *ChemElectroChem* 2022;9:e202200479. <https://doi.org/10.1002/CELC.202200479>.
- [358] Tesch MF, Neugebauer S, Narangoda P V., Schlögl R, Mechler AK. The rotating disc electrode: measurement protocols and reproducibility in the evaluation of catalysts for the oxygen evolution reaction. *Energy Advances* 2023;2:1823–30. <https://doi.org/10.1039/D3YA00340J>.
- [359] Pedro Aguiar dos Santos J, Cesar Rufino F, Yutaka Ota JI, Fernandes RC, Vicentini R, Pagan CJB, et al. Best practices for electrochemical characterization of supercapacitors. *Journal of Energy Chemistry* 2023;80:265–83. <https://doi.org/10.1016/J.JECHEM.2022.12.034>.
- [360] Ge Y, Xie X, Roscher J, Holze R, Qu Q. How to measure and report the capacity of electrochemical double layers, supercapacitors, and their electrode materials. *Journal of Solid State Electrochemistry* 2020;24:3215–30. <https://doi.org/10.1007/S10008-020-04804-X>.
- [361] Zhou S. Advanced Diagnostics of Polymer Electrolyte Fuel Cells. Doctoral Thesis, UCL (University College London) 2024.
- [362] Moffat RJ. Describing the uncertainties in experimental results. *Exp Therm Fluid Sci* 1988;1:3–17. [https://doi.org/10.1016/0894-1777\(88\)90043-X](https://doi.org/10.1016/0894-1777(88)90043-X).
- [363] Gittins JW, Chen Y, Arnold S, Augustyn V, Balducci A, Brousse T, et al. Interlaboratory study assessing the analysis of supercapacitor electrochemistry data. *J Power Sources* 2023;585. <https://doi.org/10.1016/j.jpowsour.2023.233637>.
- [364] Al-Youbi AO, Gómez de la Fuente JL, Pérez-Alonso FJ, Obaid AY, Fierro JLG, Peña MA, et al. Effects of multiwalled carbon nanotube morphology on the synthesis and electrocatalytic performance of Pt supported by multiwalled carbon nanotubes. *Appl Catal B* 2014;150–151:21–9. <https://doi.org/10.1016/J.APCATB.2013.11.043>.
- [365] Jaouen F, Herranz J, Lefèvre M, Dodelet JP, Kramm UI, Herrmann I, et al. Cross-laboratory experimental study of non-noble-metal electrocatalysts for the oxygen reduction reaction. *ACS Appl Mater Interfaces* 2009;1:1623–39. <https://doi.org/10.1021/AM900219G>.

- [366] Mamtani K, Bruening C, Co A, Ozkan US. A comparison of oxygen reduction reaction (ORR) performance for iron-nitrogen-carbon (FeNC) catalysts in acidic and alkaline media. *Res Rev Electrochem* 2017;8:107.
- [367] Noriega-Perez D. Machine Learning to support the optimization of graphene-based materials for electrochemical devices 2025. <https://doi.org/10.5281/zenodo.17038945>.
- [368] Compton OC, Nguyen ST. Graphene oxide, highly reduced graphene oxide, and graphene: Versatile building blocks for carbon-based materials. *Small* 2010;6:711–23. <https://doi.org/10.1002/smll.200901934>.
- [369] Becerril HA, Mao J, Liu Z, Stoltenberg RM, Bao Z, Chen Y. Evaluation of solution-processed reduced graphene oxide films as transparent conductors. *ACS Nano* 2008;2:463–70. <https://doi.org/10.1021/nn700375n>.
- [370] Pei S, Cheng H-M. The reduction of graphene oxide. *Carbon N Y* 2012;50:3210–28. <https://doi.org/10.1016/j.carbon.2011.11.010>.
- [371] Dreyer DR, Park S, Bielawski CW, Ruoff RS. The chemistry of graphene oxide. *Chem Soc Rev* 2010;39:228–40. <https://doi.org/10.1039/b917103g>.
- [372] De Silva KKH, Huang H-H, Joshi R, Yoshimura M. Restoration of the graphitic structure by defect repair during the thermal reduction of graphene oxide. *Carbon N Y* 2020;166:74–90. <https://doi.org/10.1016/j.carbon.2020.05.015>.
- [373] Gao X, Jang J, Nagase S. Hydrazine and thermal reduction of graphene oxide: Reaction mechanisms, product structures, and reaction design. *Journal of Physical Chemistry C* 2010;114:832–42. <https://doi.org/10.1021/jp909284g>.
- [374] Acik M, Lee G, Mattevi C, Pirkle A, Wallace RM, Chhowalla M, et al. The role of oxygen during thermal reduction of graphene oxide studied by infrared absorption spectroscopy. *Journal of Physical Chemistry C* 2011;115:19761–81. <https://doi.org/10.1021/jp2052618>.
- [375] Dolbin AV, Khlistyuck MV, Esel'Son VB, Gavrillko VG, Vinnikov NA, Basnukaeva RM, et al. The effect of the thermal reduction temperature on the structure and sorption capacity of reduced graphene oxide materials. *Appl Surf Sci* 2016;361:213–20. <https://doi.org/10.1016/j.apsusc.2015.11.167>.
- [376] Jakhar R, Yap JE, Joshi R. Microwave reduction of graphene oxide. *Carbon N Y* 2020;170:277–93. <https://doi.org/10.1016/j.carbon.2020.08.034>.
- [377] Voiry D, Yang J, Kupferberg J, Fullon R, Lee C, Jeong HY, et al. High-quality graphene via microwave reduction of solution-exfoliated graphene oxide. *Science (1979)* 2016;353:1413–6. <https://doi.org/10.1126/science.aah3398>.
- [378] Chouhan RS, Pandey A, Qureshi A, Ozguz V, Niazi JH. Nanomaterial resistant microorganism mediated reduction of graphene oxide. *Colloids Surf B Biointerfaces* 2016;146:39–46. <https://doi.org/10.1016/j.colsurfb.2016.05.053>.
- [379] Kuila T, Bose S, Khanra P, Mishra AK, Kim NH, Lee JH. A green approach for the reduction of graphene oxide by wild carrot root. *Carbon N Y* 2012;50:914–21. <https://doi.org/10.1016/j.carbon.2011.09.053>.

- [380] Chua CK, Pumera M. Chemical reduction of graphene oxide: A synthetic chemistry viewpoint. *Chem Soc Rev* 2014;43:291–312. <https://doi.org/10.1039/c3cs60303b>.
- [381] De Silva KKH, Huang H-H, Joshi RK, Yoshimura M. Chemical reduction of graphene oxide using green reductants. *Carbon N Y* 2017;119:190–9. <https://doi.org/10.1016/j.carbon.2017.04.025>.
- [382] Chen W, Yan L, Bangal PR. Chemical reduction of graphene oxide to graphene by sulfur-containing compounds. *Journal of Physical Chemistry C* 2010;114:19885–90. <https://doi.org/10.1021/jp107131v>.
- [383] Chudziak T, Montes-García V, Czepa W, Pakulski D, Musiał A, Valentini C, et al. A comparative investigation of the chemical reduction of graphene oxide for electrical engineering applications. *Nanoscale* 2023;15:17765–75. <https://doi.org/10.1039/d3nr04521h>.
- [384] Chen W, Yan L. Preparation of graphene by a low-temperature thermal reduction at atmosphere pressure. *Nanoscale* 2010;2:559–63. <https://doi.org/10.1039/b9nr00191c>.
- [385] European Union. European Chemicals Agency. <https://EchaEuropaEu> 2025.
- [386] Wan W, Zhang F, Yu S, Zhang R, Zhou Y. Hydrothermal formation of graphene aerogel for oil sorption: The role of reducing agent, reaction time and temperature. *New Journal of Chemistry* 2016;40:3040–6. <https://doi.org/10.1039/c5nj03086b>.
- [387] Baker M, Penny D. Is there a reproducibility crisis? *Nature* 2016;533:452–4. <https://doi.org/10.1038/533452A>.
- [388] Stark PB. Before reproducibility must come preproducibility world-view. *Nature* 2018;557:613. <https://doi.org/10.1038/D41586-018-05256-0>.
- [389] Fanelli D. Is science really facing a reproducibility crisis, and do we need it to? *Proc Natl Acad Sci U S A* 2018;115:2628–31. <https://doi.org/https://doi.org/10.1073/PNAS.1708272114>.
- [390] Zuliani JE, Caguiat JN, Kirk DW, Jia CQ. Considerations for consistent characterization of electrochemical double-layer capacitor performance. *J Power Sources* 2015;290:136–43. <https://doi.org/10.1016/j.jpowsour.2015.04.019>.
- [391] Noori A, El-Kady MF, Rahmanifar MS, Kaner RB, Mousavi MF. Towards establishing standard performance metrics for batteries, supercapacitors and beyond. *Chem Soc Rev* 2019;48:1272–341. <https://doi.org/10.1039/c8cs00581h>.
- [392] Mathis TS, Kurra N, Wang X, Pinto D, Simon P, Gogotsi Y. Energy Storage Data Reporting in Perspective—Guidelines for Interpreting the Performance of Electrochemical Energy Storage Systems. *Adv Energy Mater* 2019;9. <https://doi.org/10.1002/aenm.201902007>.
- [393] Balducci A, Belanger D, Brousse T, Long JW, Sugimoto W. A guideline for reporting performance metrics with electrochemical capacitors: From electrode materials to full devices. *J Electrochem Soc* 2017;164:A1487–8. <https://doi.org/10.1149/2.0851707jes>.
- [394] Zhang S, Pan N. Supercapacitors performance evaluation. *Adv Energy Mater* 2015;5. <https://doi.org/10.1002/aenm.201401401>.

- [395] Zhao J, Burke AF. Electrochemical Capacitors: Performance Metrics and Evaluation by Testing and Analysis. *Adv Energy Mater* 2021;11:2002192. <https://doi.org/10.1002/AENM.202002192>.
- [396] Frith JT, Lacey MJ, Ulissi U. A non-academic perspective on the future of lithium-based batteries. *Nature Communications* 2023 14:1 2023;14:1–17. <https://doi.org/10.1038/s41467-023-35933-2>.
- [397] Basha N, Arcucci R, Angeli P, Anastasiou C, Abadie T, Casas CQ, et al. Machine learning and physics-driven modelling and simulation of multiphase systems. *International Journal of Multiphase Flow* 2024;179. <https://doi.org/10.1016/j.ijmultiphaseflow.2024.104936>.
- [398] Iannello S, Friso A, Galvanin F, Materazzi M. A Hybrid Physics-Machine Learning Approach for Modeling Plastic-Bed Interactions during Fluidized Bed Pyrolysis. *Energy and Fuels* 2025;39:4549–64. <https://doi.org/10.1021/acs.energyfuels.4c05870>.
- [399] Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nature Reviews Physics* 2021;3:422–40. <https://doi.org/10.1038/s42254-021-00314-5>.
- [400] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* 2019;378:686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- [401] Reinbold PAK, Kageorge LM, Schatz MF, Grigoriev RO. Robust learning from noisy, incomplete, high-dimensional experimental data via physically constrained symbolic regression. *Nat Commun* 2021;12. <https://doi.org/10.1038/s41467-021-23479-0>.
- [402] Petsagkourakis P, Galvanin F. Safe model-based design of experiments using Gaussian processes. *Comput Chem Eng* 2021;151. <https://doi.org/10.1016/j.compchemeng.2021.107339>.
- [403] Liang Q, Gongora AE, Ren Z, Tiihonen A, Liu Z, Sun S, et al. Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains. *NPJ Comput Mater* 2021;7. <https://doi.org/10.1038/s41524-021-00656-9>.
- [404] Zhang Y, Apley DW, Chen W. Bayesian Optimization for Materials Design with Mixed Quantitative and Qualitative Variables. *Sci Rep* 2020;10. <https://doi.org/10.1038/s41598-020-60652-9>.
- [405] Agarwal G, Doan HA, Robertson LA, Zhang L, Assary RS. Discovery of Energy Storage Molecular Materials Using Quantum Chemistry-Guided Multiobjective Bayesian Optimization. *Chemistry of Materials* 2021;33:8133–44. <https://doi.org/10.1021/acs.chemmater.1c02040>.
- [406] Griffiths R-R, Klärner L, Moss H, Ravuri A, Truong S, Stanton S, et al. GAUCHE: A Library for Gaussian Processes in Chemistry. *Adv Neural Inf Process Syst*, vol. 36, 2023.
- [407] Deisenroth MP, Luo Y, van der Wilk M. A Practical Guide to Gaussian Processes. <https://Infallible-Thompson-49de36NetlifyApp/> 2020.

- [408] Tfaily A, Kokkolaras M, Bartoli N, Diouane Y. Efficient Acquisition Functions for Bayesian Optimization in the Presence of Hidden Constraints. *AIAA Aviation and Aeronautics Forum and Exposition, AIAA AVIATION Forum 2023* 2023. <https://doi.org/10.2514/6.2023-4261>.
- [409] Alazmi A, El Tall O, Rasul S, Hedhili MN, Patole SP, Costa PMFJ. A process to enhance the specific surface area and capacitance of hydrothermally reduced graphene oxide. *Nanoscale* 2016;8:17782–7. <https://doi.org/10.1039/C6NR04426C>.
- [410] Wang Z, Tang Z, Han Z, Shen S, Zhao B, Yang J. Effect of drying conditions on the structure of three-dimensional N-doped graphene and its electrochemical performance. *RSC Adv* 2015;5:19838–43. <https://doi.org/10.1039/c4ra15494k>.
- [411] Gadipelli S, Akbari H, Li J, Howard CA, Zhang H, Shearing PR, et al. Structure-guided Capacitance Relationships in Oxidized Graphene Porous Materials Based Supercapacitors. *Energy and Environmental Materials* 2023;6. <https://doi.org/10.1002/eem2.12637>.
- [412] Gadipelli S, Guo J, Li Z, Howard CA, Liang Y, Zhang H, et al. Understanding and Optimizing Capacitance Performance in Reduced Graphene-Oxide Based Supercapacitors. *Small Methods* 2023;7:2201557. <https://doi.org/10.1002/smt.202201557>.
- [413] Tamboli SH, Seok Kim B, Choi G, Lee H, Lee D, Patil UM, et al. Post-heating effects on the physical and electrochemical capacitive properties of reduced graphene oxide paper. *J Mater Chem A* 2014;2:5077–86. <https://doi.org/10.1039/c4ta00209a>.
- [414] Mei BA, Munteshari O, Lau J, Dunn B, Pilon L. Physical Interpretations of Nyquist Plots for EDLC Electrodes and Devices. *Journal of Physical Chemistry C* 2018;122:194–206. <https://doi.org/10.1021/ACS.JPCC.7B10582>.
- [415] Gadipelli S, Akbari H, Li J, Howard CA, Zhang H, Shearing PR, et al. Structure-guided Capacitance Relationships in Oxidized Graphene Porous Materials Based Supercapacitors. *Energy & Environmental Materials* 2023;6:e12637. <https://doi.org/10.1002/EEM2.12637>.
- [416] Pedro Aguiar dos Santos J, Cesar Rufino F, Yutaka Ota JI, Fernandes RC, Vicentini R, Pagan CJB, et al. Best practices for electrochemical characterization of supercapacitors. *Journal of Energy Chemistry* 2023;80:265–83. <https://doi.org/10.1016/j.jechem.2022.12.034>.
- [417] Su H, Lin S, Deng S, Lian C, Shang Y, Liu H. Predicting the capacitance of carbon-based electric double layer capacitors by machine learning. *Nanoscale Adv* 2019;1:2162–6. <https://doi.org/10.1039/C9NA00105K>.
- [418] Ha T, Kim SK, Choi JW, Chang H, Jang HD. pH controlled synthesis of porous graphene sphere and application to supercapacitors. *Advanced Powder Technology* 2019;30:18–22. <https://doi.org/10.1016/J.APT.2018.10.002>.
- [419] Géron A. Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: Concepts, tools, and techniques to build intelligent systems. “ O’Reilly Media, Inc.”; 2022.

- [420] Ding R, Chen J, Chen Y, Liu J, Bando Y, Wang X. Unlocking the potential: machine learning applications in electrocatalyst design for electrochemical hydrogen energy transformation. *Chem Soc Rev* 2024. <https://doi.org/10.1039/d4cs00844h>.
- [421] Suvarna M, Pérez-Ramírez J. Embracing data science in catalysis research. *Nat Catal* 2024;7:624–35. <https://doi.org/10.1038/s41929-024-01150-3>.
- [422] Yao Z, Lum Y, Johnston A, Mejia-Mendoza LM, Zhou X, Wen Y, et al. Machine learning for a sustainable energy future. *Nat Rev Mater* 2023;8:202–15. <https://doi.org/10.1038/s41578-022-00490-5>.
- [423] Lai NS, Tew YS, Zhong X, Yin J, Li J, Yan B, et al. Artificial Intelligence (AI) Workflow for Catalyst Design and Optimization. *Ind Eng Chem Res* 2023;62:17835–48. <https://doi.org/10.1021/acs.iecr.3c02520>.
- [424] McCullough K, Williams T, Mingle K, Jamshidi P, Lauterbach J. High-throughput experimentation meets artificial intelligence: A new pathway to catalyst discovery. *Physical Chemistry Chemical Physics* 2020;22:11174–96. <https://doi.org/10.1039/d0cp00972e>.
- [425] Nandy A, Duan C, Taylor MG, Liu F, Steeves AH, Kulik HJ. Computational Discovery of Transition-metal Complexes: From High-throughput Screening to Machine Learning. *Chem Rev* 2021;121:9927–10000. <https://doi.org/10.1021/acs.chemrev.1c00347>.
- [426] Lombardo T, Duquesnoy M, El-Bouysidy H, Àren F, Gallo-Bueno A, Jørgensen PB, et al. Artificial Intelligence Applied to Battery Research: Hype or Reality? *Chem Rev* 2022;122:10899–969. <https://doi.org/10.1021/acs.chemrev.1c00108>.
- [427] Toyao T, Maeno Z, Takakusagi S, Kamachi T, Takigawa I, Shimizu KI. Machine Learning for Catalysis Informatics: Recent Applications and Prospects. *ACS Catal* 2020;10:2260–97. <https://doi.org/10.1021/acscatal.9b04186>.
- [428] MacLeod BP, Parlane FGL, Rupnow CC, Dettelbach KE, Elliott MS, Morrissey TD, et al. A self-driving laboratory advances the Pareto front for material properties. *Nat Commun* 2022;13. <https://doi.org/10.1038/s41467-022-28580-6>.
- [429] Dave A, Mitchell J, Kandasamy K, Wang H, Burke S, Paria B, et al. Autonomous Discovery of Battery Electrolytes with Robotic Experimentation and Machine Learning. *Cell Rep Phys Sci* 2020;1. <https://doi.org/10.1016/j.xcrp.2020.100264>.
- [430] Han Z, Chen A, Li Z, Zhang M, Wang Z, Yang L, et al. Machine learning-based design of electrocatalytic materials towards high-energy lithium||sulfur batteries development. *Nat Commun* 2024;15:8433. <https://doi.org/10.1038/s41467-024-52550-9>.
- [431] Merchant A, Batzner S, Schoenholz SS, Aykol M, Cheon G, Cubuk ED. Scaling deep learning for materials discovery. *Nature* 2023;624:80–5. <https://doi.org/10.1038/s41586-023-06735-9>.
- [432] Zhong M, Tran K, Min Y, Wang C, Wang Z, Dinh CT, et al. Accelerated discovery of CO₂ electrocatalysts using active machine learning. *Nature* 2020;581:178–83. <https://doi.org/10.1038/s41586-020-2242-8>.

- [433] Moliner M, Román-Leshkov Y, Corma A. Machine Learning Applied to Zeolite Synthesis: The Missing Link for Realizing High-Throughput Discovery. *Acc Chem Res* 2019;52:2971–80. <https://doi.org/10.1021/acs.accounts.9b00399>.
- [434] Gundry L, Guo S-X, Kennedy G, Keith J, Robinson M, Gavaghan D, et al. Recent advances and future perspectives for automated parameterisation, Bayesian inference and machine learning in voltammetry. *Chemical Communications* 2021;57:1855–70. <https://doi.org/10.1039/D0CC07549C>.
- [435] Ciucci F, Chen C. Analysis of electrochemical impedance spectroscopy data using the distribution of relaxation times: A Bayesian and hierarchical Bayesian approach. *Electrochim Acta* 2015;167:439–54. <https://doi.org/10.1016/j.electacta.2015.03.123>.
- [436] Gundry L, Kennedy G, Bond AM, Zhang J. Establishing zone regions in cyclic voltammetry using unsupervised machine learning. *Journal of Electroanalytical Chemistry* 2023;942. <https://doi.org/10.1016/j.jelechem.2023.117551>.
- [437] Liu J, Ciucci F. The Gaussian process distribution of relaxation times: A machine learning tool for the analysis and prediction of electrochemical impedance spectroscopy data. *Electrochim Acta* 2020;331. <https://doi.org/10.1016/j.electacta.2019.135316>.
- [438] Karan V, Gallant MC, Fei Y, Ceder G, Persson KA. Ion correlations explain kinetic selectivity in diffusion-limited solid state synthesis reactions. 2025.
- [439] Pillai HS, Li Y, Wang SH, Omidvar N, Mu Q, Achenie LEK, et al. Interpretable design of Ir-free trimetallic electrocatalysts for ammonia oxidation with graph neural networks. *Nat Commun* 2023;14. <https://doi.org/10.1038/s41467-023-36322-5>.
- [440] Zheng H, Sivonxay E, Christensen R, Gallant M, Luo Z, McDermott M, et al. The ab initio non-crystalline structure database: empowering machine learning to decode diffusivity. *NPJ Comput Mater* 2024;10:295. <https://doi.org/10.1038/s41524-024-01469-2>.
- [441] Goodall REA, Lee AA. Predicting materials properties without crystal structure: deep representation learning from stoichiometry. *Nat Commun* 2020;11. <https://doi.org/10.1038/s41467-020-19964-7>.
- [442] Noda K, Wakiuchi A, Hayashi Y, Yoshida R. Advancing extrapolative predictions of material properties through learning to learn using extrapolative episodic training. *Commun Mater* 2025;6. <https://doi.org/10.1038/s43246-025-00754-x>.
- [443] Sendek AD, Yang Q, Cubuk ED, Duerloo KAN, Cui Y, Reed EJ. Holistic computational structure screening of more than 12 000 candidates for solid lithium-ion conductor materials. *Energy Environ Sci* 2017;10:306–20. <https://doi.org/10.1039/c6ee02697d>.
- [444] Artrith N, Lin Z, Chen JG. Predicting the activity and selectivity of bimetallic metal catalysts for ethanol reforming using machine learning. *ACS Catal* 2020;10:9438–44. <https://doi.org/10.1021/acscatal.0c02089>.
- [445] Zhou S, Wu Y, Xu L, Kockelmann W, Rasha L, Du W, et al. Water content estimation in polymer electrolyte fuel cells using synchronous electrochemical

- impedance spectroscopy and neutron imaging. *Cell Rep Phys Sci* 2024;5:102208. <https://doi.org/10.1016/j.xcrp.2024.102208>.
- [446] Galiounas E, Tranter TG, Owen RE, Robinson JB, Shearing PR, Brett DJL. Battery state-of-charge estimation using machine learning analysis of ultrasonic signatures. *Energy and AI* 2022;10:100188. <https://doi.org/10.1016/J.EGYAI.2022.100188>.
- [447] Fordham A, Milojevic Z, Giles E, Du W, Owen RE, Michalik S, et al. Correlative non-destructive techniques to investigate aging and orientation effects in automotive Li-ion pouch cells. *Joule* 2023;7:2622–52. <https://doi.org/10.1016/j.joule.2023.10.011>.
- [448] Galiounas E, Owen RE, Robinson JB, Jervis R. The generalisation challenge: Assessment of the efficacy of acoustic signals for state estimation of lithium-ion batteries via machine learning. *J Power Sources* 2025;630:236047. <https://doi.org/10.1016/j.jpowsour.2024.236047>.
- [449] Fordham A, Joo S-B, Owen RE, Galiounas E, Buckwell M, Brett DJL, et al. Investigating the Performance and Safety of Li-Ion Cylindrical Cells Using Acoustic Emission and Machine Learning Analysis. *J Electrochem Soc* 2024;171:070521. <https://doi.org/10.1149/1945-7111/ad59c9>.
- [450] Galiounas E, Iacoviello F, Mirza M, Rasha L, Owen RE, Robinson JB, et al. Investigations into the Dynamic Acoustic Response of Lithium-Ion Batteries During Lifetime Testing. *J Electrochem Soc* 2024;171:070514. <https://doi.org/10.1149/1945-7111/ad5d21>.
- [451] Petrich L, Westhoff D, Feinauer J, Finegan DP, Daemi SR, Shearing PR, et al. Crack detection in lithium-ion cells using machine learning. *Comput Mater Sci* 2017;136:297–305. <https://doi.org/10.1016/j.commatsci.2017.05.012>.
- [452] Bailey JJ, Chen J, Hack J, Perez-Page M, Holmes SM, Brett DJL, et al. Lab-based X-ray micro-computed tomography coupled with machine-learning segmentation to investigate phosphoric acid leaching in high-temperature polymer electrolyte fuel cells. *J Power Sources* 2021;509:230347. <https://doi.org/10.1016/j.jpowsour.2021.230347>.
- [453] Bailey JJ, Wade A, Boyce AM, Zhang YS, Brett DJL, Shearing PR. Quantitative assessment of machine-learning segmentation of battery electrode materials for active material quantification. *J Power Sources* 2023;557:232503. <https://doi.org/10.1016/j.jpowsour.2022.232503>.
- [454] Zhou S, Shearing PR, Brett DJL, Jervis R. Machine learning as an online diagnostic tool for proton exchange membrane fuel cells. *Curr Opin Electrochem* 2022;31:100867. <https://doi.org/10.1016/j.coelec.2021.100867>.
- [455] Zhang T, Tan R, Zhu P, Zhang TY, Huang J. Unlocking Ultrafast Diagnosis of Retired Batteries via Interpretable Machine Learning and Optical Fiber Sensors. *ACS Energy Lett* 2025;10:862–71. <https://doi.org/10.1021/acseenergylett.4c03054>.
- [456] Cai J, Chu X, Xu K, Li H, Wei J. Machine learning-driven new material discovery. *Nanoscale Adv* 2020;2:3115–30. <https://doi.org/10.1039/d0na00388c>.

- [457] Oviedo F, Ferres JL, Buonassisi T, Butler KT. Interpretable and Explainable Machine Learning for Materials Science and Chemistry. *Acc Mater Res* 2022;3:597–607. <https://doi.org/10.1021/accountsmr.1c00244>.
- [458] Schrier J, Norquist AJ, Buonassisi T, Brgoch J. In Pursuit of the Exceptional: Research Directions for Machine Learning in Chemical and Materials Science. *J Am Chem Soc* 2023;145:21699–716. <https://doi.org/10.1021/jacs.3c04783>.
- [459] Zhong X, Gallagher B, Liu S, Kailkhura B, Hiszpanski A, Han TYJ. Explainable machine learning in materials science. *NPJ Comput Mater* 2022;8. <https://doi.org/10.1038/s41524-022-00884-7>.
- [460] Ghiringhelli LM, Vybiral J, Levchenko S V., Draxl C, Scheffler M. Big data of materials science: Critical role of the descriptor. *Phys Rev Lett* 2015;114:105503. <https://doi.org/10.1103/PHYSREVLETT.114.105503>.
- [461] Krenn M, Pollice R, Guo SY, Aldeghi M, Cervera-Lierta A, Friederich P, et al. On scientific understanding with artificial intelligence. *Nature Reviews Physics* 2022;4:761–9. <https://doi.org/10.1038/s42254-022-00518-3>.
- [462] Carleo G, Cirac I, Cranmer K, Daudet L, Schuld M, Tishby N, et al. Machine learning and the physical sciences. *Rev Mod Phys* 2019;91. <https://doi.org/10.1103/RevModPhys.91.045002>.
- [463] Bessa MA, Bostanabad R, Liu Z, Hu A, Apley DW, Brinson C, et al. A framework for data-driven analysis of materials under uncertainty: Countering the curse of dimensionality. *Comput Methods Appl Mech Eng* 2017;320:633–67. <https://doi.org/10.1016/j.cma.2017.03.037>.
- [464] Bianco A, Cheng H-M, Enoki T, Gogotsi Y, Hurt RH, Koratkar N, et al. All in the graphene family - A recommended nomenclature for two-dimensional carbon materials. *Carbon* N Y 2013;65:1–6. <https://doi.org/10.1016/j.carbon.2013.08.038>.
- [465] Akima H. A new method of interpolation and smooth curve fitting based on local procedures. *Journal of the ACM (JACM)* 1970;17:589–602.
- [466] Barrett EP, Joyner LG, Halenda PP. The determination of pore volume and area distributions in porous substances. I. Computations from nitrogen isotherms. *J Am Chem Soc* 1951;73:373–80.
- [467] Abadi M, Barham P, Chen J, Chen Z, Davis A, Dean J, et al. TensorFlow: A system for large-scale machine learning. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016*, 2016, p. 265–83.
- [468] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 2011;12:2825–30.
- [469] Chawla A. Why R-Squared is a flawed regression metric. <https://www.dailydoseofds.com/why-r-squared-is-a-flawed-regression-metric/> 2023.
- [470] Akaike H. A New Look at the Statistical Model Identification. *IEEE Trans Automat Contr* 1974;19:716–23. <https://doi.org/10.1109/TAC.1974.1100705>.

- [471] Schwarz G. Estimating the dimension of a model. *The Annals of Statistics* 1978;461–4.
- [472] Lundberg SM, Allen PG, Lee S-I. A Unified Approach to Interpreting Model Predictions. *Adv Neural Inf Process Syst* 2017;30.
- [473] Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, et al. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence* 2020 2:1 2020;2:56–67. <https://doi.org/10.1038/s42256-019-0138-9>.
- [474] Śliwak A, Grzyb B, Díez N, Gryglewicz G. Nitrogen-doped reduced graphene oxide as electrode material for high rate supercapacitors. *Appl Surf Sci* 2017;399:265–71. <https://doi.org/10.1016/j.apsusc.2016.12.060>.
- [475] Li SM, Yang SY, Wang YS, Tsai HP, Tien HW, Hsiao ST, et al. N-doped structures and surface functional groups of reduced graphene oxide and their effect on the electrochemical performance of supercapacitor with organic electrolyte. *J Power Sources* 2015;278:218–29. <https://doi.org/10.1016/j.jpowsour.2014.12.025>.
- [476] Zhang LL, Gu Y, Zhao XS. Advanced porous carbon electrodes for electrochemical capacitors. *J Mater Chem A Mater* 2013;1:9395–408. <https://doi.org/10.1039/c3ta11114h>.
- [477] Chabi S, Peng C, Hu D, Zhu Y. Ideal three-dimensional electrode structures for electrochemical energy storage. *Advanced Materials* 2014;26:2440–5. <https://doi.org/10.1002/adma.201305095>.
- [478] Zhang F, Liu T, Hou G, Kou T, Yue L, Guan R, et al. Hierarchically porous carbon foams for electric double layer capacitors. *Nano Res* 2016;9:2875–88. <https://doi.org/10.1007/s12274-016-1173-z>.
- [479] Yao B, Chandrasekaran S, Zhang H, Ma A, Kang J, Zhang L, et al. 3D-Printed Structure Boosts the Kinetics and Intrinsic Capacitance of Pseudocapacitive Graphene Aerogels. *Advanced Materials* 2020;32. <https://doi.org/10.1002/adma.201906652>.
- [480] Kim J, Eum JH, Kang J, Kwon O, Kim H, Kim DW. Tuning the hierarchical pore structure of graphene oxide through dual thermal activation for high-performance supercapacitor. *Sci Rep* 2021;11. <https://doi.org/10.1038/s41598-021-81759-7>.
- [481] Li Z, Gadipelli S, Yang Y, He G, Guo J, Li J, et al. Exceptional supercapacitor performance from optimized oxidation of graphene-oxide. *Energy Storage Mater* 2019;17:12–21. <https://doi.org/10.1016/j.ensm.2018.12.006>.
- [482] Li Z, Gadipelli S, Li H, Howard CA, Brett DJL, Shearing PR, et al. Tuning the interlayer spacing of graphene laminate films for efficient pore utilization towards compact capacitive energy storage. *Nat Energy* 2020;5:160–8. <https://doi.org/10.1038/s41560-020-0560-6>.
- [483] Li Z, Gadipelli S, Yang Y, Guo Z. Design of 3D Graphene-Oxide Spheres and Their Derived Hierarchical Porous Structures for High Performance Supercapacitors. *Small* 2017;13. <https://doi.org/10.1002/smll.201702474>.

- [484] Yu M, Xie B, Yang Y, Zhang Y, Chen Y, Yu W, et al. Residual oxygen groups in nitrogen-doped graphene to enhance the capacitive performance. *RSC Adv* 2017;7:15293–301. <https://doi.org/10.1039/c6ra28596a>.
- [485] Morimoto N, Kubo T, Nishina Y. Tailoring the oxygen content of graphite and reduced graphene oxide for specific applications. *Sci Rep* 2016;6. <https://doi.org/10.1038/srep21715>.
- [486] Gadipelli S, Lu Y, Skipper NT, Yildirim T, Guo Z. Design of hyperporous graphene networks and their application in solid-amine based carbon capture systems. *J Mater Chem A Mater* 2017;5:17833–40. <https://doi.org/10.1039/c7ta05789j>.
- [487] Yu X, Park HS. Sulfur-incorporated, porous graphene films for high performance flexible electrochemical capacitors. *Carbon N Y* 2014;77:59–65. <https://doi.org/10.1016/j.carbon.2014.05.002>.
- [488] Chen X, Chen X, Xu X, Yang Z, Liu Z, Zhang L, et al. Sulfur-doped porous reduced graphene oxide hollow nanosphere frameworks as metal-free electrocatalysts for oxygen reduction reaction and as supercapacitor electrode materials. *Nanoscale* 2014;6:13740–7. <https://doi.org/10.1039/c4nr04783d>.
- [489] Yu X, Kang Y, Park HS. Sulfur and phosphorus co-doping of hierarchically porous graphene aerogels for enhancing supercapacitor performance. *Carbon N Y* 2016;101:49–56. <https://doi.org/10.1016/j.carbon.2016.01.073>.
- [490] Tarimo DJ, Oyedotun KO, Mirghni AA, Manyala N. Sulphur-reduced graphene oxide composite with improved electrochemical performance for supercapacitor applications. *Int J Hydrogen Energy* 2020;45:13189–201. <https://doi.org/10.1016/j.ijhydene.2020.03.059>.
- [491] Molnar C. Interpretable Machine Learning A Guide for Making Black Box Models Explainable. GitHub 2025. <https://christophm.github.io/interpretable-ml-book/> (accessed January 2, 2026).
- [492] Wager S, Hastie T, Efron B. Confidence Intervals for Random Forests: The Jackknife and the Infinitesimal Jackknife. vol. 15. 2014.
- [493] Zhao X, Dong H, Xiao Y, Hu H, Cai Y, Liang Y, et al. Three-dimensional Nitrogen-doped graphene as binder-free electrode materials for supercapacitors with high volumetric capacitance and the synergistic effect between nitrogen configuration and supercapacitive performance. *Electrochim Acta* 2016;218:32–40. <https://doi.org/10.1016/j.electacta.2016.09.096>.
- [494] Karaman C, Bayram E, Karaman O, Aktaş Z. Preparation of high surface area nitrogen doped graphene for the assessment of morphologic properties and nitrogen content impacts on supercapacitors. *Journal of Electroanalytical Chemistry* 2020;868. <https://doi.org/10.1016/j.jelechem.2020.114197>.
- [495] Dang F, Zhao W, Yang P, Wu H, Liu Y. Nitrogen and sulfur co-doped hierarchical graphene hydrogel for high-performance electrode materials. *J Appl Electrochem* 2020;50:463–73. <https://doi.org/10.1007/s10800-020-01404-5>.
- [496] Nankya R, Opar DO, Kim MJ, Paek SM, Jung H. Synergetic effect of nitrogen and sulfur co-doping in mesoporous graphene for enhanced energy storage properties in supercapacitors and lithium-ion batteries. *J Solid State Chem* 2020;289. <https://doi.org/10.1016/j.jssc.2020.121451>.

- [497] Kan Y, Ning G, Ma X. Sulfur-decorated nanomesh graphene for high-performance supercapacitors. *Chinese Chemical Letters* 2017;28:2277–80. <https://doi.org/10.1016/j.ccllet.2017.11.026>.
- [498] Zhu Z, Wang Z, Ba Z, Dong J, Zhang Q, Zhao X. Three-dimensional N/S Co-doped holey graphene oxide based hydrogel electrodes for high performance supercapacitors. *J Energy Storage* 2021;39. <https://doi.org/10.1016/j.est.2021.102658>.
- [499] Pham TV, Kim JG, Jung JY, Kim JH, Cho H, Seo TH, et al. High Areal Capacitance of N-Doped Graphene Synthesized by Arc Discharge. *Adv Funct Mater* 2019;29. <https://doi.org/10.1002/adfm.201905511>.
- [500] Sun H, Wang Q, Wu T, Miao Y, Fang Y. Plasma-assisted synthesis of pyrrolic-nitrogen doped reduced graphene oxide to enhance supercapacitor performance. *Appl Surf Sci* 2020;527. <https://doi.org/10.1016/j.apsusc.2020.146574>.
- [501] Sun L, Wang L, Tian C, Tan T, Xie Y, Shi K, et al. Nitrogen-doped graphene with high nitrogen level via a one-step hydrothermal reaction of graphene oxide with urea for superior capacitive energy storage. *RSC Adv* 2012;2:4498–506. <https://doi.org/10.1039/c2ra01367c>.
- [502] Xu J, Tan Z, Zeng W, Chen G, Wu S, Zhao Y, et al. A Hierarchical Carbon Derived from Sponge-Templated Activation of Graphene Oxide for High-Performance Supercapacitor Electrodes. *Advanced Materials* 2016;28:5222–8. <https://doi.org/10.1002/adma.201600586>.
- [503] Wang H, Sun X, Liu Z, Lei Z. Creation of nanopores on graphene planes with MgO template for preparing high-performance supercapacitor electrodes. *Nanoscale* 2014;6:6577–84. <https://doi.org/10.1039/c4nr00538d>.
- [504] Li J, Liu K, Gao X, Yao B, Huo K, Cheng Y, et al. Oxygen- and Nitrogen-Enriched 3D Porous Carbon for Supercapacitors of High Volumetric Capacity. *ACS Appl Mater Interfaces* 2015;7:24622–8. <https://doi.org/10.1021/acsami.5b06698>.
- [505] Li Z, Yang B, Li K, Wang H, Lv X, Guo Y, et al. Freestanding 3D mesoporous graphene oxide for high performance energy storage applications. *RSC Adv* 2014;4:51640–7. <https://doi.org/10.1039/c4ra08519a>.
- [506] Huang Y, Gao A, Song X, Shu D, Yi F, Zhong J, et al. Supramolecule-Inspired Fabrication of Carbon Nanoparticles in Situ Anchored Graphene Nanosheets Material for High-Performance Supercapacitors. *ACS Appl Mater Interfaces* 2016;8:26775–82. <https://doi.org/10.1021/acsami.6b08511>.
- [507] Singh KP, Bhattacharjya D, Razmjooei F, Yu JS. Effect of pristine graphene incorporation on charge storage mechanism of three-dimensional graphene oxide: Superior energy and power density retention. *Sci Rep* 2016;6. <https://doi.org/10.1038/srep31555>.
- [508] Fujisawa K, Cruz-Silva R, Yang KS, Kim YA, Hayashi T, Endo M, et al. Importance of open, heteroatom-decorated edges in chemically doped-graphene for supercapacitor applications. *J Mater Chem A Mater* 2014;2:9532–40. <https://doi.org/10.1039/c4ta00936c>.

- [509] Zhang Y, Tao B, Xing W, Zhang L, Xue Q, Yan Z. Sandwich-like nitrogen-doped porous carbon/graphene nanoflakes with high-rate capacitive performance. *Nanoscale* 2016;8:7889–98. <https://doi.org/10.1039/c5nr05151g>.
- [510] Sun F, Gao J, Pi X, Wang L, Yang Y, Qu Z, et al. High performance aqueous supercapacitor based on highly nitrogen-doped carbon nanospheres with unimodal mesoporosity. *J Power Sources* 2017;337:189–96. <https://doi.org/10.1016/j.jpowsour.2016.10.086>.
- [511] Tian W, Gao Q, Zhang L, Yang C, Li Z, Tan Y, et al. Renewable graphene-like nitrogen-doped carbon nanosheets as supercapacitor electrodes with integrated high energy-power properties. *J Mater Chem A Mater* 2016;4:8690–9. <https://doi.org/10.1039/c6ta02828d>.
- [512] Islam MM, Subramaniam CM, Akhter T, Faisal SN, Minett AI, Liu HK, et al. Three dimensional cellular architecture of sulfur doped graphene: self-standing electrode for flexible supercapacitors, lithium ion and sodium ion batteries. *J Mater Chem A Mater* 2017;5:5290–302. <https://doi.org/10.1039/c6ta10933k>.
- [513] Phiri J, Gane P, Maloney TC. Multidimensional Co-Exfoliated Activated Graphene-Based Carbon Hybrid for Supercapacitor Electrode. *Energy Technology* 2019;7. <https://doi.org/10.1002/ente.201900578>.
- [514] Lin Y, Chen H, Shi Y, Wang G, Chen L, Wang F, et al. Nitrogen and Sulfur Co-Doped Graphene-Like Carbon from Industrial Dye Wastewater for Use as a High-Performance Supercapacitor Electrode. *Global Challenges* 2019;3. <https://doi.org/10.1002/gch2.201900043>.
- [515] Nanaji K, Upadhyayula V, Rao TN, Anandan S. Robust, Environmentally Benign Synthesis of Nanoporous Graphene Sheets from Biowaste for Ultrafast Supercapacitor Application. *ACS Sustain Chem Eng* 2019;7:2516–29. <https://doi.org/10.1021/acssuschemeng.8b05419>.