Automated Deep Learning in Ophthalmology: AI That Can Build AI

Ciara O'Byrne[1], Abdallah Abbas[1], Edward Korot[1,2], Pearse A Keane[1,3*]

[1]Moorfields Eye Hospital NHS Foundation Trust, London, United Kingdom

[2]Byers Eye Institute, Stanford University, California, United States

[3]NIHR Biomedical Research Centre for Ophthalmology, Moorfields Eye Hospital NHS Foundation, London, United Kingdom

* = corresponding author

Corresponding author:

Dr. Pearse A. Keane

Moorfields Eye Hospital NHS Foundation Trust, London, United Kingdom

pearsek@gmail.com

Abstract

Purpose of review: The purpose of this review is to describe the current status of automated deep learning in healthcare and to explore and detail the development of these models using commercially available platforms. We highlight key studies demonstrating the effectiveness of this technique and discuss current challenges and future directions of automated deep learning.

Recent findings: There are several commercially-available automated deep learning platforms. While specific features differ between platforms, they utilise the common approach of supervised learning. Ophthalmology is an exemplar specialty in the area, with a number of recent proof-of-concept studies exploring classification of retinal fundus photographs, optical coherence tomography images and indocyanine green angiography images. Automated deep learning has also demonstrated impressive results in other specialties such as dermatology, radiology and histopathology.

Summary: Automated deep learning allows users without coding expertise to develop deep learning algorithms. It is rapidly establishing itself as a valuable tool for those with limited technical experience. Despite residual challenges, it offers considerable potential in the future of patient management, clinical research and medical education.

Keywords: Automated deep learning, artificial medical intelligence, code-free deep learning, deep learning

Introduction

Aging populations, changing disease patterns and the rise in patient autonomy and expectations are contributing to unprecedented pressures on our healthcare systems (1–3). Rising levels of clinician burn-out (4), enormous administrative burdens and strained resources additionally intensify the situation resulting in a healthcare ecosystem that is unsustainable and not fit for purpose. Deep learning, a subtype of artificial intelligence (AI) inspired by the neural architecture of the human brain, has risen to prominence as a potential solution to some of these challenges. In fact, a report from 2020 predicted that the Global AI in Healthcare Market would grow from USD 4.9 billion in 2020 to USD 45.2 billion by 2026 (5). However, despite considerable promise, deep learning is limited in healthcare by the need for highly specialised technical expertise, advanced computing resources and significant financial investment.

AI That Can Build AI

More recently, automated deep learning has emerged showing promising results across a number of areas (6–8). Described by the New York Times as “AI That Can Build AI" (9), it is a technique that automates the process of pre-processing, network architecture selection, and hyperparameter tuning allowing those without coding expertise to develop models. It is commercially available on a number of different platforms including Amazon Rekognition Custom Labels (Amazon), Apple Create ML (Apple), Baidu EasyDL (Baidu), Clarifai Train (Clarifai), Google Cloud AutoML Vision (Google), Huawei ModelArts ExeML (Huawei), MedicMind Deep learning Training Platform (MedicMind), and Microsoft Azure Custom Vision (Microsoft). Automated deep learning has sparked considerable excitement within the fields of healthcare and research by offering to obviate the barriers that have traditionally limited the accessibility of deep learning. With its heavy use of imaging, ophthalmology is particularly suited to applications of deep learning. To date, it has been one of the leading specialties in the exploration of this technique (6,10,11). By enabling clinicians to create their own deep learning models, automated deep learning may truly maximise the way in which we harness the power of data and AI leading to novel discoveries, applications and improvements in patient care.

The automated deep learning process

Automated *machine* learning (AutoML) describes a set of techniques that *assist* with dataset management, model selection, and optimisation of hyperparameters. Methods include Auto-WEKA, Auto-Sklearn, TuPAQ and AlphaD3M (12). Although AutoML automates part of the machine learning pipeline, it still requires coding expertise (13). Automated *deep* learning is based on an approach called neural architecture search and typically utilises reinforcement learning algorithms to automatically develop a deep learning architecture. Reinforcement learning describes a goal-oriented reward process in which the algorithm learns through trial and error. While specific features vary between platforms, the basic principles for the development of an automated deep learning model are similar. The graphical user interfaces (GUI) are intuitive and offer drag-and-drop or simple upload tools, removing the need for any coding expertise. Furthermore, cloud-based approaches have removed the need for large local computing power. For each of the commercial platforms examined by the authors, the GUI consists of three common components - data upload, data visualization, and model evaluation. Most platforms (Amazon, Apple, Clarifai, Google and Microsoft) offer the option to develop models for image classification, segmentation and object detection, with Google additionally providing the facility to create models using tabular data. We will now describe the development process of an image classifier model in detail (see figure 1).

Image Classification

Automated deep learning image classifiers use supervised learning, a machine learning technique in which the model discovers patterns from labelled input data and adjusts its internal parameters to output a prediction algorithm with the lowest possible error rate (14).

Automated deep learning does not remove the initial task of data preparation. This is a critical stage and the clinician must be cognisant to the importance of well-labelled datasets that are representative of the use-case and target population. It is also imperative that data ethics and governance are adhered to if public datasets are not being used. Therefore, dataset curation and labeling represents a persistent pain-point within automated deep learning. A number of companies have begun to release services to address this challenge including Amazon Automate Data Labeling, Clarifai Scribe Label and Google Cloud AutoML Vision Human Labeling.

Once the dataset has been curated, the project can be created and named directly via the GUI. For those using Amazon, Clarifai, MedicMind and Microsoft, the dataset should be pre-organised into labelled folders before upload. Amazon additionally allows for the project to be linked to a cloud bucket of labelled images. Google offers the user the option to either upload the dataset directly from the computer using the GUI or to convert it into .csv files locally and upload via a cloud storage bucket, this is useful for managing large datasets and labelsets. MedicMind also allows for .csv files to be used, however, not through a cloud bucket. Apple does not use cloud computing, and labels are assigned according to local folders. After the dataset has been uploaded, the labels can be reviewed and amended if necessary. Metrics are also supplied alerting the user to the distribution of images per label. The dataset can then be either manually or automatically split into three parts. The training set receives approximately 60-80% of the images and is important for network parameter selection. The remainder are divided between a validation set, used to optimise the model parameters, and a held-out independent test set, which ultimately assesses the model performance.

Once the user is satisfied with the uploaded dataset, the model may be trained. Following the training process, detailed statistics for the model performance are provided, which vary between platforms. Confidence threshold is provided by Amazon, Clarifai, Google and

Microsoft. This can be altered to generate new precision and recalls by all of the above except for Amazon. Confusion matrices (Apple, Clarifai, Google and MedicMind) allow the user to visualise the true positives and false negatives and are essential for model evaluation. Precision recall curves are provided by Amazon, Clarifai and Google. MedicMind is the only automated deep learning platform examined that offers saliency maps, an approach being explored within the sub-field of explainable AI (15). Although both MedicMind and Google offer external validation, Google is the only platform familiar to the authors that allows for external validation via batch prediction, allowing predictions to be generated on a large external dataset efficiently. Download of the final model is facilitated by Google and Microsoft platforms.

Automated deep learning in the literature

Ophthalmology

In 2019, our group published one of the earliest demonstrations of automated deep learning for medical imaging classification (6). Two clinicians with no coding expertise developed automated deep learning models using five publicly available datasets of retinal fundus images, optical coherence tomography (OCT) images, dermatological skin lesion images and chest x-ray images. The models were developed using Google AutoML Vision. Sensitivity (recall), specificity, positive predictive value (precision) and area under the precision recall curve (AUPRC) were used to evaluate model discriminative performance and diagnostic properties. Aside from the multi-label model trained using one of the chest x-ray datasets, we were able to demonstrate comparable accuracy to state-of-the-art bespoke deep learning systems. Similarly, Kim and colleagues used Google AutoML Vision to train two models in the classification of pachychoroid disease (11). A dataset of 783 ultra-widefield indocyanine green angiography (ICG) images was curated, labelled by two retina specialists. The model performance was assessed using precision and recall, with accuracy levels then compared against both ophthalmic residents and retinal specialists. The authors reported that their second model demonstrated better precision and accuracy than the retina specialists with comparable recall and specificity. In comparison to the ophthalmic residents, the second model demonstrated inferior recall and specificity but greater precision and accuracy.

More recently, our group published a comprehensive performance and featureset review of six commercially available platforms using four open-source ophthalmic imaging datasets, including two retinal fundus photograph datasets and two OCT datasets (10). Twenty four automated deep learning models were trained by clinicians with none to limited coding expertise, and the specific features and performance of each application programming interface were evaluated. Notably, only Amazon, Apple, Google and Microsoft had the ability to process large imaging datasets and of these, Apple's performance was considerably worse than Amazon. We postulate that this may be due to Apple Create ML running locally, rather than utilising large cloud computing resources. We also observed an improved performance with OCT classification models across all platforms in comparison to the fundus photograph models. We suspect this may be due to the increased dimensionality of colour fundus photographs. As we have previously highlighted, Google AutoML Vision is the only commercial deep learning platform allowing the user to carry out external validation via batch prediction. The caveat is that this must be carried out using the command line interface, thus requiring some degree of coding experience (16).

Other specialties

Automated deep learning has also been applied in a number of other specialties. As discussed, the automated deep learning model we described in 2019 demonstrated impressive results in the classification of chest x-rays, with one model showing comparable performance to bespoke deep learning models (6). More recently, chest x-ray classification has been explored by other research groups using Microsoft Custom Vision (7) and Google AutoML Vision (17). Google AutoML Vision has also been used to develop image classifier models in histopathology (18), neuro-histopathology (8) and otolaryngology (19) while Wang et al. utilised the Google AutoML object detection tool to develop a system capable of identifying and risk stratifying high risk mutations in thyroid nodules (20). Borkowski and colleagues performed a comparison between Google AutoML Vision versus Apple Create ML in a variety of different lung and colon diagnostic pathology scenarios (21). The authors trained twelve deep learning models in total (six on each platform) to differentiate between a variety of lung and colon pathologies. While the authors did not determine any statistically significant differences in terms of model performance between both platforms, they observe that although Apple Create ML models are limited to the local computer, Google AutoML Vision utilises Google Cloud resulting in computing fees.

Tabular data

Current automated deep learning-based ophthalmology research has focused on interpreting fundus photographs, ICG angiography and OCT scans (10,11, 22–24). Structured data, based on a tabular format of columns and rows, represents an additional rich source of information relating to patient histories, diagnoses and prognoses. The potential benefit of such data within ophthalmology research is exemplified by projects such as the Intelligent Research in Sight (IRIS®) Registry. This contains information from nearly 66 million patients (25). Thus, the diversification of automated deep learning-based ophthalmology research to include models that take advantage of structured inputs represents a significant step forward.

Though the current literature is scarce, initial models built using structured datasets have shown promise. A recent study by Antaki et al. demonstrated that ophthalmologists with no programming experience could use electronic health record data to build predictive models for proliferative vitreoretinopathy, using an interactive application in MATLAB (26). These code-free models achieved comparable F1 scores to manually coded models built on the same datasets. Moreover, novel tools specifically engineered for structured data have now been developed, such as Google Cloud's AutoML Tables. This platform enables clinicians to build classification, regression and time-series machine learning models without needing to code. Early work using this platform includes a model that predicts visual outcomes in patients receiving treatment for neovascular age-related macular degeneration. This achieved an area under the receiver operating characteristic curve (AUROC) of 0.892 (27). To assist in scheduling, our group has trained a cataract surgery time prediction model which predicts operating time with a mean absolute error of 5 minutes. Future work should aim to further examine the feasibility of such tools in comparison to conventional machine learning methods, emulating the numerous comparative studies exploring automated deep learning for image classification (6).

Limitations of automated deep learning

Automated deep learning is not a panacea. Though outside the scope of this article, there are barriers common to all AI applications in healthcare. These include dataset curation, ethical and medicolegal considerations, data governance and regulatory issues as well as patient and clinician acceptance of these systems.

The ‘black box’ phenomena is well-documented as a limitation in the implementation of artificial medical intelligence tools (28–30). This is further intensified by the inability to select or obtain information about the neural architecture framework chosen for the model. Given the fact that minimal technical expertise is required for the development of these automated systems, it is imperative that robust tools are developed to allow the clinician to understand how the model has reached its decision. Discriminatory bias is another issue that must be highlighted to clinicians with limited deep learning experience when developing these models. Discriminatory bias describes the situation in which a model is selected to optimally represent the majority population. This may result in an inferior performance with under-represented groups. Although public datasets represent a valuable resource, they may be particularly prone to discriminatory bias depending on how the dataset was collected and who the deep learning model is being developed for. Clinicians must be aware of the perils associated with overfitting and develop models with their target population in mind. External validation, with datasets representing various real-world image acquisition environments with varying patient demographics, remains key.

There are also limitations specific to automated deep learning platforms. These platforms do not offer flexibility in selecting between models architectures used to train the model. The evaluation metrics vary between platforms, which can make it difficult to accurately assess and compare the performance of models *between* platforms. Many platforms do not offer the facility to externally validate the model. This is an essential step in the process, and one which must be incorporated into a system if it is to be considered for implementation. Finally, there are costs associated with these commercial platforms, particularly those that are cloud-based. Although automated deep learning is an important step towards the democratisation of AI in healthcare, it still presents financial burdens which may be challenging for small research groups with minimal funding.

Future directions

Direct patient care

Automated deep learning has the potential to play an important role in patient care, particularly as effectiveness improves and ethical and governance regulations are established. Clinicians are use-case experts, who are best suited to train models, specified for patient-relevant endpoints. Consequently, clinicians are best suited to apply the relevant labels for model training. By allowing physicians to independently devise and develop deep learning models, patient needs may be uniquely and efficiently addressed. Image recognition models may greatly enhance screening programmes, particularly in under-resourced areas (31). Structured data approaches may prove useful in the prediction of patient outcomes, while natural language processing may alleviate the significant administrative burdens clinicians are faced with at present. Despite these advantages, hospital management and clinicians must be aware that the use of such models for direct patient care would also be subject to the same clinical validation and regulatory requirements as bespoke deep learning systems.

Clinical research

Automated deep learning may radically enhance the clinical research landscape. With the capacity to play a number of different roles within the research toolkit, it has the potential to both alleviate the strain of laborious administrative tasks while also identifying new patterns within data previously unknown to humans, leading the way towards clinical trial selection, drug discovery and development. To first ascertain if there is sufficient signal to invest in further custom model development via coding, automated deep learning models may be trained as a proof of concept.

Improved technology

While cloud computing has alleviated some of the challenges associated with deep learning, it still depends on high bandwidth, low latency and robust privacy safeguards (32). Further advances in mobile technology, such as 5G, may address these issues through the use of automated deep learning systems via local edge models (i.e., compact low power models which do not require a continuous internet connection to run) (33,34). Combined with telemedicine and wearable sensors, these models may considerably improve the quality of healthcare in under-resourced communities.

Medical education

It is essential that medical students and clinicians alike are adequately equipped with the skills needed to navigate the field of artificial medical intelligence (35). Automated deep learning may be able to assist with this in a number of different ways (see figure 2). First, it provides a practical opportunity to grasp the process of model development, evaluation and implementation. Second, it draws attention to the potential hazards and limitations associated with deep learning models and in turn, imparts a deeper understanding of the ethical considerations surrounding these systems. As discussed, automated deep learning allows the medical community to develop algorithms that are uniquely appropriate for their specific needs. This could be capitalised upon to develop automated deep learning models to enhance medical education with applications in surgical training, disease recognition and monitoring of progress and performance.

Conclusion

Automated deep learning is establishing itself as a potential solution to many of the challenges facing healthcare systems today. We believe it will play a central role in the future democratisation and industrialisation of artificial intelligence in healthcare, ultimately transforming patient care, medical education and clinical research.

Key points

- Automated deep learning allows users with no coding expertise to develop deep learning algorithms.
- It has shown impressive results in comparison to bespoke deep learning models across a number of specialties, including ophthalmology, histopathology, and radiology.
- It is available on a number of commercial platforms and, with technological advances, is likely to become even more accessible in the future.
- Automated deep learning represents a promising tool in the future of patient care, clinical research and medical education.

Acknowledgements

Dr. Korot is supported by a Springboard Grant from the Moorfields Eye Charity.

Dr Keane is supported by a Moorfields Eye Charity Career Development Award (R190028A) and a UK Research & Innovation Future Leaders Fellowship (MR/T019050/1).

Disclosures

Dr. Keane has acted as a consultant for DeepMind, Roche, Novartis, Apellis, and BitFount and is an equity owner in Big Picture Medical. He has received speaker fees from Heidelberg Engineering, Topcon, Allergan, and Bayer. Dr. Korot has acted as a consultant for Google Health. He has received consulting fees from Genentech.

References and recommended reading

Papers of particular significance, published within the annual period of review, have been highlighted :

- * for special interest
- ** for outstanding interest

1. GBD 2019 Viewpoint Collaborators. Five insights from the Global Burden of Disease Study 2019. Lancet. 2020 Oct 17;396(10258):1135–59.

2. Seniori Costantini A, Gallo F, Pega F, Saracci R, Veerus P, West R. Population health and status of epidemiology in Western European, Balkan and Baltic countries. Int J Epidemiol. 2015 Feb;44(1):300–23.

3. Maresova P, Javanmardi E, Barakovic S, Barakovic Husic J, Tomsone S, Krejcar O, et al. Consequences of chronic diseases and other limitations associated with old age - a scoping review. BMC Public Health. 2019 Nov 1;19(1):1431.

4. The Lancet. Physician burnout: a global crisis. Lancet. 2019 Jul 13;394(10193):93.

5. Artificial Intelligence in Healthcare Market with Covid-19 Impact Analysis by Offering (Hardware, Software, Services), Technology (Machine Learning, NLP, Context-Aware Computing, Computer Vision), End-Use Application, End User and Region - Global Forecast to 2026. Markets and Markets; 2020 Jun. Report No.: 5116503.

6. **Faes L, Wagner SK, Fu DJ, Liu X, Korot E, Ledsam JR, et al. Automated deep learning design for medical image classification by health-care professionals with no coding experience: a feasibility study. Lancet Digit Health. 2019 Sep;1(5):e232–42. This is a seminal publication in the field of automated deep learning for medical image classification. The authors demonstrate that two clinicians with no coding expertise can develop deep learning models using automated deep learning via Google Cloud AutoML Vision that have comparable performance to bespoke deep learning models.

7. *Borkowski AA, Viswanadhan NA, Thomas LB, Guzman RD, Deland LA, Mastorides SM. Using Artificial Intelligence for COVID-19 Chest X-ray Diagnosis. Fed Pract. 2020 Sep;37(9):398–404. This paper discusses the development of an automated deep learning system to identify COVID-19 from a public dataset of chest x-ray images. It is one of the few publications in the field of automated deep learning for medical image classification using Microsoft CustomVision.

8. Koga S, Ghayal NB, Dickson DW. Deep Learning-Based Image Classification in Differentiating Tufted Astrocytes, Astrocytic Plaques, and Neuritic Plaques. J Neuropathol Exp Neurol. 2021 Mar 22;80(4):306–12.

9. Metz C. Building AI that can build AI. The New York Times Çevrimiçi (Erişim, 4 Şubat 2018): https://www nytimes com/2017/11/05/technolo gy/machine-learning-artificial-intelligence-ai html. 2017;

10. **Korot E, Guan Z, Ferraz D, Wagner SK, Zhang G, Liu X, et al. Code-free Deep

Learning for Multi-Modality Medical Image Classification. Nature Machine Intelligence. This is the only comprehensive review of six commercially available automated deep learning platforms that we are aware of. This article details the various features of each platform and describes the performance accuracy in comparison to previously published bespoke models.

11. * Kim IK, Lee K, Park JH, Baek J, Lee WK. Classification of pachychoroid disease on ultrawide-field indocyanine green angiography using auto-machine learning platform. Br J Ophthalmol [Internet]. 2020 Jul 3; Available from: http://dx.doi.org/10.1136/bjophthalmol-2020-316108 This is the first application of automated deep learning to indocyanine green angiography images. It highlights the effectiveness of automated deep learning and discusses some limitations.

12. Waring J, Lindvall C, Umeton R. Automated machine learning: Review of the state-of-the-art and opportunities for healthcare. Artif Intell Med. 2020 Apr;104:101822.

13. Quanming Yao, Mengshuo Wang, Yuqiang Chen, Wenyuan Dai, Yu-Feng Li, Wei-Wei Tu, Qiang Yang, Yang Yu. Taking Human Out of Learning Applications: A Survey on Automated Machine Learning. arXiv:181013306v4 [Internet]. 2019 Dec 16; Available from: https://arxiv.org/abs/1810.13306

14. LeCun Y, Bengio Y, Hinton G. Deep learning. Nature. 2015 May 28;521(7553):436–44.

15. Adebayo J, Gilmer J, Muelly M, Goodfellow I, Hardt M, Kim B. Sanity Checks for Saliency Maps [Internet]. arXiv [cs.CV]. 2018. Available from: http://arxiv.org/abs/1810.03292

16. Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, et al. Identifying Medical Diagnoses and Treatable Diseases by Image-Based Deep Learning. Cell. 2018 Feb 22;172(5):1122–31.e9.

17. Ghosh T, Tanwar S, Chumber S, Vani K. Classification of chest radiographs using

general purpose cloud-based automated machine learning: pilot study. Egyptian Journal of Radiology and Nuclear Medicine. 2021 May 5;52(1):120.

18. Zeng Y, Zhang J. A machine learning model for detecting invasive ductal carcinoma with Google Cloud AutoML Vision. Comput Biol Med. 2020 Jul;122:103861.

19. Livingstone D, Chau J. Otoscopic diagnosis using computer vision: An automated machine learning approach. Laryngoscope. 2020 Jun;130(6):1408–13.

20. Wang S, Xu J, Tahmasebi A, Daniels K, Liu J-B, Curry J, et al. Incorporation of a Machine Learning Algorithm With Object Detection Within the Thyroid Imaging Reporting and Data System Improves the Diagnosis of Genetic Risk. Front Oncol. 2020 Nov 12;10:591846.

21. Borkowski AA, Wilson CP, Borkowski SA, Brannon Thomas L, Deland LA, Grewe SJ, et al. Google Auto ML versus Apple Create ML for Histopathologic Cancer Diagnosis; Which Algorithms Are Better? [Internet]. arXiv [q-bio.QM]. 2019. Available from: http://arxiv.org/abs/1903.08057

22. Korot E, Wagner S, Faes L, Fu DJ, Liu X, Ferraz D, et al. AI building AI: Deep Learning Detection of Referable Diabetic Retinopathy Sans-coding. Invest Ophthalmol Vis Sci. 2020 Jun 10;61(7):2025–2025.

23. Beqiri S, Abbas A, Korot E, Wagner S, Struyven R, Kelly M. Investigating the impact of saliency maps on clinician’s confidence in model predictions. In.

24. Wagner S, Korot E, Khalid H, Heeren T, Ferraz D, Guan Z, et al. Automated Machine Learning Model for Fundus Photo Gradeability and Laterality: A Public ML Research Toolkit Sans-coding. Invest Ophthalmol Vis Sci. 2020 Jun 10;61(7):2029–2029.

25. Larkin H. Iris Registry update [Internet]. EuroTimes. 2021 [cited 2021 May 3]. Available from: https://www.eurotimes.org/iris-registry-update/

26. Antaki F, Kahwati G, Sebag J, Coussa RG, Fanous A, Duval R, et al. Predictive modeling of proliferative vitreoretinopathy using automated machine learning by ophthalmologists without coding experience. Sci Rep. 2020 Nov 11;10(1):19528.

27. Abbas A, Beqiri S, Wagner S, Korot E, Singh R, Struyven R. Using the What-if Tool to perform nearest counterfactual analysis on an AutoML model that predicts visual acuity outcomes in patients receiving treatment for wet age-related macular degeneration. In.

28. The Lancet Respiratory Medicine. Opening the black box of machine learning. Lancet Respir Med. 2018 Nov;6(11):801.

29. Zihni E, Madai VI, Livne M, Galinovic I, Khalil AA, Fiebach JB, et al. Opening the black box of artificial intelligence for clinical decision support: A study predicting stroke outcome. PLoS One. 2020 Apr 6;15(4):e0231166.

30. London AJ. Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability [Internet]. Vol. 49, Hastings Center Report. 2019. p. 15–21. Available from: http://dx.doi.org/10.1002/hast.973

31. Bellemo V, Lim ZW, Lim G, Nguyen QD, Xie Y, Yip MYT, et al. Artificial intelligence using deep learning to screen for referable and vision-threatening diabetic retinopathy in Africa: a clinical validation study. Lancet Digit Health. 2019 May;1(1):e35–44.

32. Merenda M, Porcaro C, Iero D. Edge Machine Learning for AI-Enabled IoT Devices: A Review. Sensors [Internet]. 2020 Apr 29;20(9). Available from: http://dx.doi.org/10.3390/s20092533

33. Keane PA, Topol EJ. Medicine and meteorology: Cloud, connectivity, and care. Lancet. 2020 Apr 25;395(10233):1334.

34. Greco L, Percannella G, Ritrovato P, Tortorella F, Vento M. Trends in IoT based solutions for health care: Moving AI to the edge. Pattern Recognit Lett. 2020

Jul;135:346–53.

35. Keane PA, Topol EJ. AI-facilitated health care requires education of clinicians. Lancet. 2021 Apr 3;397(10281):1254.

Figures

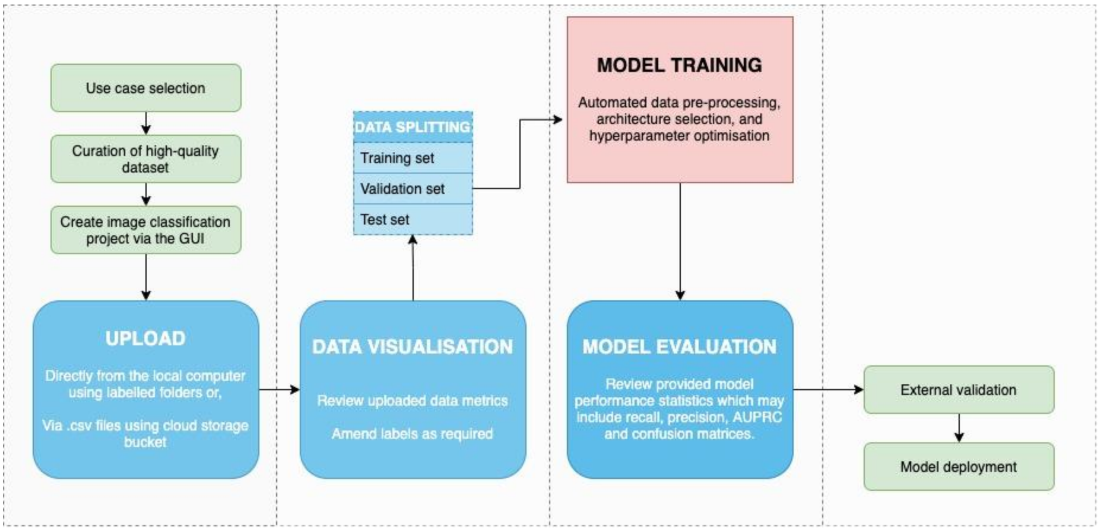

Figure 1: The development process of an automated deep learning image classifier model.

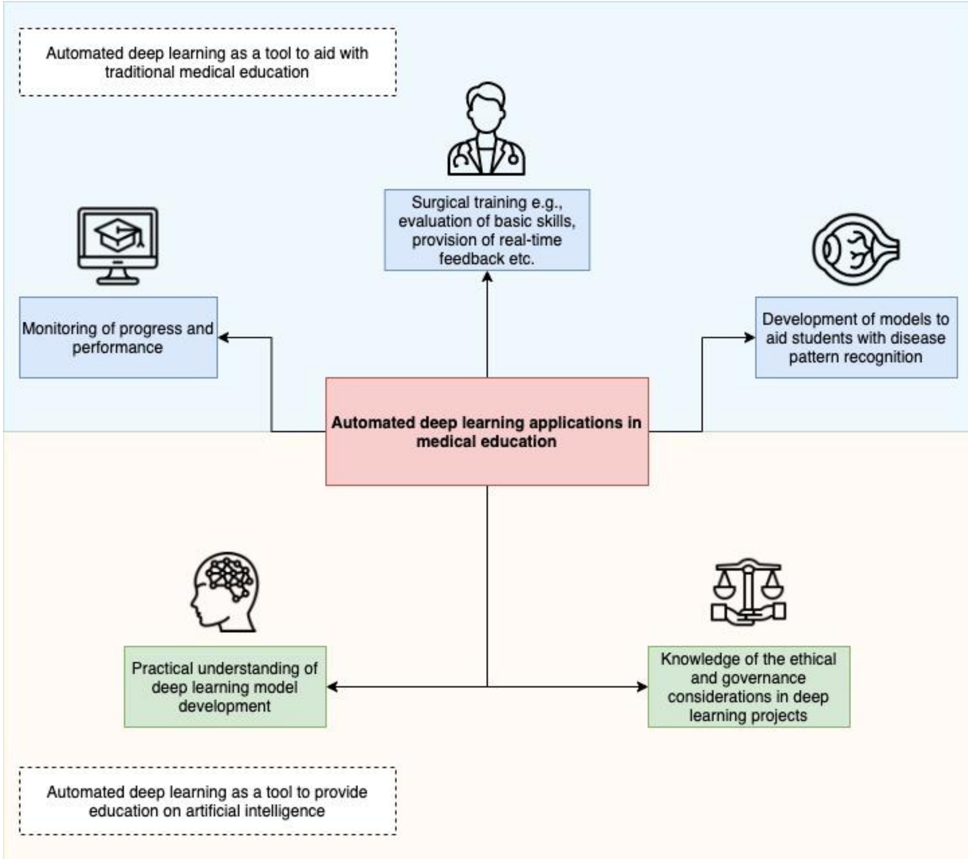

Figure 2: Overview of the potential applications of automated deep learning in medical education.