Chapter 1. Introduction to computational biomedicine

Shunzhou Wan[1], Peter V. Coveney[1,2,3]*

[1]Centre for Computational Science, Department of Chemistry, University College London

[2]Advanced Research Computing Centre, University College London

[3]Computational Science Laboratory, Institute for Informatics, Faculty of Science, University of Amsterdam

Author Information

*Corresponding Author, E-mail: p.v.coveney@ucl.ac.uk

Keywords: molecular modelling, machine learning, binding affinity, computer aided drug design, clinical decision support systems, digital twin

ABSTRACT

The domain of computational biomedicine is a new and burgeoning one. Its areas of concern cover all scales of human biology, physiology and pathology, commonly referred to as medicine, from the genomic to the whole human and beyond, including epidemiology and population health. Computational biomedicine aims to provide high fidelity descriptions and predictions of the behaviour of biomedical systems of both fundamental scientific and clinical importance. Digital twins and virtual humans aim to reproduce the extremely accurate duplicate of real-world human beings in cyber space, which can be used to make highly accurate predictions that take complicated conditions into account. When that can be done reliably enough for the predictions to be actionable, such an approach will make an impact in the pharmaceutical industry by reducing or even replacing the extremely laboratory intensive preclinical process of making and testing compounds in laboratories, and in clinical applications by assisting clinicians to make diagnostic and treatment decisions.

1. Introduction

Technologies originating with the digital revolution over the past twenty years have dramatically evolved multiple paradigms in the biomedicine and healthcare sectors. The use of computer models and simulation is now widespread to accelerate drug development in the pharmaceutical industry and to aid in diagnosis and treatment of diseases in clinical application. One of the major advantages of computational modelling is that it provides insight into the physical, chemical and biological bases of clinical interventions. These could be the underlying molecular interactions and mechanisms of drugs to their target biomolecules or a virtual three-dimensional view of the region of a patient's anatomy on which surgeons may operate. Molecular level insights are often inaccessible experimentally while details of organs may not be visible without invasive procedures.

As these computational models become more reliable, one would hope to use these methods to quantitatively predict the outcome of experiments or clinical operations prior to performing them. The computational methods may come to replace many experimental procedures. In this way, computational techniques should reduce time and cost in the industrial processes of drug discovery, which takes on average of more than 10 years and $2.6 billion to bring a new drug to market [1]. As far as clinicians are concerned, they can invoke clinical decision support systems (CDSS) to enhance decision-making in the clinical workflow. CDSS encompass a variety of tools, some based on knowledge, others on models and computation. While the use of knowledge management is still the main feature of CDSS, the components based on modelling and computation are starting to make an impact on decision making [2]. A recent study showed that a personalised CDSS may be created by constructing a computational pipeline at the molecular level with individual patients' genomic data to select the optimal drugs for a given patient, based on the predicted ranking of the binding efficacy of the particular drugs to variant proteins [3].

Computational modelling can also convert images obtained from medical scans into highly realistic virtual 3D models which can be used for training and for rehearsing complex cases

prior to surgery. Due to these potential benefits, computer-based techniques have been adopted as routine techniques by the scientific community, along with the pharmaceutical industry and the healthcare sector.

The relentless enhancement in the performance of high-end computers, especially with advances in emerging embedded and many-core GPU accelerators of increasing diversity, is another key factor accounting for the increasing adoption of computer-based methods in science and industry in general, and in biomedicine in particular, over recent decades. In 1977, a "supercomputer" of that time—an IBM System/370 Model 168—was used for the first molecular dynamics simulation of a small protein [4] which had a top speed of a million (10^6) floating point operations per second (FLOPS), or a MegaFLOPS. In June 2022, the first machine to officially reach the exascale was announced, which achieved a peak performance of 1.102 exaFLOPS, or 1.102 quintillion (i.e. a billion billion or 10^{18}) FLOPS. The machine is Frontier, built at Oak Ridge National Laboratory (ORNL) in Tennessee, USA. Just one second of job execution on the entirety of Frontier would have needed more than 31,000 years to complete for a "supercomputer" from 45 years ago.

This increase in computing power enables us to perform complex calculations at very high speeds, with the simulation settings closer to real world situations. The first MD simulation of a protein mentioned above [4] simulated a small protein *in vacuo* for 9.2 picoseconds. Today it is routine to run simulations for systems consisting of tens to hundreds of thousands of atoms for tens of nanoseconds or longer. During the COVID-19 pandemic, computational modelling and simulation play an important role in the understanding of the SARS-CoV-2 virus and in the early drug development to find potential therapeutic compounds. One of the studies, which won the ACM Gordon Bell award in 2020, used Summit supercomputer to simulate the spike protein and viral envelope with a model consisting of 305 million atoms [5]. As one of the authors stated: "we are giving people never-before-seen, intimate views of this virus, with resolution that is impossible to achieve experimentally". The molecular systems we study today are much more realistic as to their physiological conditions, with proteins being

surrounded by water and ions, frequently including portions of cells membranes if required. The force fields used to describe the interactions between all the atoms in these simulations have also significantly improved, helping in some respects to make the simulations more accurate.

It should be noted that, while we can run such spatially enormous structures, it does not mean we can run them for long enough to extract anything useful from them. In the particular case of spike protein and viral envelope simulation [5], a total of ~702 ns trajectories were collected. The biological processes such as receptor binding and membrane fusion, however, occur over much longer timescales, from microseconds to hours. The larger the simulation domain, the longer the timescale one must run over; owing to the serial nature of time evolution of almost all codes employed on supercomputers, this unfortunately often means that the processes of interest get further and further out of reach for the computers available. This is why we need to use multiscale methods, to bridge between time scales. During the COVID-19 pandemic, for example, multiple mathematical models were used to evaluate the virus spread, to control of infectious diseases, and to define an optimal strategy for vaccine administration [6].

The speed with which such simulations are executed is itself of critical concern as, if it can be done quickly enough, enables decision-makers to take appropriate actions in a restricted “time-critical” window of time. Pharmaceutical companies need to quickly weed out compounds that may have toxicity issues or poor pharmacokinetics in preclinical studies, the so-called “fail fast, fail early” strategy, before making expensive late-stage investments during clinical trials. It is hoped that such work can be accomplished *in silico* rather than actual experimental assays. In healthcare, a computerised CDSS need to deliver actionable recommendations to clinicians at the point of care, on time scales that are sufficiently rapid to be used in a decision-making context. The ability for clinicians to determine correct patient-specific interventions quickly should improve the efficiency of healthcare, providing better patient outcomes while eliminating unnecessary costs.

2. Methods and protocols

Computational approaches have been widely used in biomedicine. Here we will summarise a few areas where the impacts are representative.

2.1. Drug development

A major step in drug discovery is to identify lead compounds which have high binding affinities with a protein target. This happens at the early preclinical stage. The binding affinity, also known as the free energy of binding, is the single most important initial indicator of drug potency, and the most challenging to predict. It can be determined experimentally by a number of methods, with various accuracies but is generally expensive, error prone and time consuming. Alternatively, one may seek to calculate the binding energy computationally. Here, methods drawn from computational chemistry offer a route forward; these are primarily based on *in silico* molecular dynamics (MD), for which several approaches to determining the free energy are possible.

The most common approaches for free energy predictions are the endpoint and the alchemical approaches [7]. As the name indicates, the endpoint methods are based on simulations of the final physical states of a system, namely the bound and unbound states in the drug binding case. The most well-known endpoint methods are molecular mechanics Poisson–Boltzmann surface area (MM/PBSA) and molecular mechanics generalized Born surface area (MM/GBSA) [8]. The alchemical approach computer the free energy change by a “alchemical” path which is nonphysical. The process transforms, for example, one ligand into another, or a given amino acid residue into a mutated form, by morphing one chemical species of group directly into another. As the free energy is a state function, it does not matter which path, physical or nonphysical, is used to compute changes in it. MD simulations are performed using alchemical transformations to accelerate otherwise much more expensive and less reliable physical pathways to compute free energies. To further improve the precision and to

generate reproducible predictions, ensemble-based approaches have been development, both for the endpoint and alchemical methods. The formal is called *enhanced sampling of molecular dynamics with approximation of continuum solvent* (ESMACS) [9], and the latter called *thermodynamic integration with enhanced sampling* (TIES) [10].

The use of endpoint and alchemical approaches is not mutually exclusive but indeed can be even more powerful when performed in tandem. The former is less accurate but also less computationally expensive than the latter, which is well suited for use in the initial hit-to-lead activities within drug discovery. The alchemical approaches, on the other hand, are most relevant to lead optimization following the identification of promising lead compounds. A combination of the two methods, along with machine learning approaches, in a workflow has been proposed to accelerate COVID-19 drug discovery on high performance computers [11]. While the currently ongoing COVID-19 pandemic has a negative impact in many areas, the computer-aided drug discovery market does experience a boost, in which the aforementioned approaches have been extensively applied, separately or jointly, to find novel drug candidates and to reposition existing drugs. An exciting global collaboration, called the COVID Moonshot, has come together for the discovery of new, urgent drug treatment for COVID-19 [12]. We ourselves are currently participating in a large-scale collaboration in which ML, docking, endpoint and alchemical approaches are applied interactively to find promising drug candidates from libraries consisting of billions of compounds. The most attractive drug candidates are subsequently being studied experimentally with some under consideration for inclusion in possible clinical trials.

2.2. Personalised medicine

In the post-genomic era, we are all very familiar with the notion that one-size-fits-all has limited validity in medicine. A drug that works well in one person might not be effective for another at all. While percentage of patients responding fall in the range of 50–75% for many drugs, it can be as low as 25% for cancer chemotherapy [13]. This is because the way that a particular

medication works depends on how it interacts with other molecules in our bodies. A small variation, one amino acid mutation in one protein, for example, can determine a drug as an effective medication or a medication that doesn't work. In the past decade, the field of oncology has witnessed substantial changes in the way patients with cancer are managed, with increasing focus on personalised medicine based on genomic variants of individual patients. Pharmacogenomics, a combination of pharmacology and genomics, provides invaluable information on how an individual's genetic profile influences the response to medication. It enables personalised medicine which uses therapeutics to a subset of patients based on their specific genetic and molecular feature, who are expected to benefit. Two mutations in the kinase domain of epidermal growth factor receptor (EGFR), exon 19 deletions and L858R substitution in exon 21, have been commonly reported in patients with non-small cell lung carcinoma (NSCLC) tumours. The kinase inhibitor drugs Gefitinib and Erlotinib are effective in patients with these mutations, while inefficient for patients without. Even having these mutations, drug resistances can arise through other mutations such as the so-called gatekeeper mutation T790M.

One main reason for the occurrence of drug resistance arises from the binding affinity changes caused by mutations in the primary sequence of amino acid residues within the protein target. We have shown in our own previous publications how the ensemble-based binding free energy methods can be used to assess functional and mechanistic impacts of mutations in the case of EGFR [14], FGFR1 (fibroblast growth factor receptor 1) [15] and ER (estrogen receptor) [3] variants. The provision of extensive genomic sequencing technologies and the rising number of large-scale tumour molecular profiling programs across institutions worldwide have revolutionized the field of precision oncology. The treatment of patients with breast cancer, for example, has shifted from the standard of care, in which all patients receive similar interventions (such as mastectomy, axillary dissection, radiotherapy and chemotherapy), to personalized medicine, where molecular characteristics of individual patients is used for prognostics, risk assessment and selection of medical interventions [16]. Realistically,

however, we still have a long way to go to deliver patient specific treatments comprehensively across the healthcare system, whereas the stratification of the population into clusters/groups which have closer similarities might well help get finer grained differentiation so as to provide better treatments to these distinct groups. As a potential component in CDSS [2], the computation of binding free energies can be used to provide recommendations for more accurate and personalised treatment based on patients' specific genomic variances. In the longer term, a related approach could be used to design new drugs which are resistant to such mutations (see subsection 2.1. Drug development).

2.3. Medical diagnosis with machine learning

While artificial intelligence (AI) is the broad science of computer algorithms designed to mimick human cognitive functions such as learning and problem-solving, machine learning is a specific subset of AI that uses datasets to train a computer to learn observable phenomena [17]. Machine learning is particularly powerful for image-based pattern recognition and has become integral to feature detection in biomedical imaging. This is frequently of considerable benefit in the initial stages of segmentation and reconstruction of complex three-dimensional geometries. Machine learning also has considerable promise in classifying categories of observed behaviour and as a less computationally demanding surrogate for inclusion within clinically based decision support systems. While the availability of diagnostic images is rapidly increasing, they are underutilized in many countries and regions due to the lack of trained diagnostic specialists. Machine learning therefore brings considerable promise to clinical practice with medical images. A survey of the diagnostic performance based on medical imaging shows that deep learning models perform on a par with healthcare professionals [18]. By adding causal reasoning into machine leaning, the prediction can achieve expert clinical accuracy, i.e. an accuracy placing in the top 25% of doctors in the cohort [19]. An AI technology, known as InnerEye, has been used in some NHS hospitals to automate lengthy radiotherapy preparations for image guided radiotherapy. The UK government has recently

set a pioneering £100m NHS consortium to use a ground-breaking artificial intelligence (AI) tool to speed up the diagnosis of breast cancer.

It should be noted, however, that the biggest drawback in AL/ML is its total lack of explanatory power. Medical decision-making can and will never be made on the basis that an AI system advised someone to take an action. Mechanistic understanding is essential. That is why, in a new book co-authored by one of us (PVC) and Roger Highfield [20], we talk about “Big AI” – meaning AI which pays attention to our scientific understanding and conforms with the laws of nature (e.g. of physics and chemistry) and understand the structural characteristics of the problems being studied. There is vast amount of biology and medicine for which mechanistic understanding is lacking (exacerbated by the perpetual pressure to “translate” any basic medical findings to the clinic as soon as possible) so there is plenty of room for improvement. AI methods are riddled with biases of many kinds which are often implicit and may easily be missed by those producing such solutions; and there are lots of open problems to face once similar data is integrated from different sources, even if nominally being from the same measurements. In short, there are often big but hidden uncertainties in these systems.

To derive the greatest benefit from working with AI, especially in healthcare, it is best to think of them like very junior colleagues. They are good at handling data and can learn quickly, but are prone to making mistakes. The indispensable expert-in-the-loop in AI-based models allows clinicians and scientists to respond to output of the software, explaining what and why, while themselves making the key clinical decisions. These responses from clinicians can in turn become training data for the AI to learn from.

2.4. Human digital twins

As defined by the National Cancer Institute at Frederick, “A digital twin is a virtual model used to make predictions and run simulations without disrupting or harming the real-world object”. It originates in engineering around 20 years ago, and is only now being seriously applied to medicine and biology. Here we focus on biomedicine, where the definition of digital twins

covers different scales, from the cellular, organ, whole human, to population levels. For the purpose of personalised medicine, the digital twin, also known by different names including digital patient, “avatar”, virtual human and so on, is built by integrating and processing vast amounts of data from individual patients to create the virtual human body. It is evolving model which keep integrating new data and predicting future statuses.

The ultimate purpose of human digital twins is to use patient-specific modelling to provide support for actionable medical and clinical decisions. High-resolution models can be evaluated virtually for many possible kinds of medical interventions in order to find the best option for the patient. This includes the optimal drugs which benefit most for the patient with minimal side effects or the optimal clinical operation strategies. Although digital twinning is still at an early stage of development and adoption in the healthcare sector, the concept is already finding some actual as well as potential application areas including patient monitoring, disease prediction, personalised treatment, population health studies as well as *in silico* clinical and other kinds of health trials. Companies like ELEMBio (https://www.elem.bio/) spun out from Barcelona Supercomputing Center (BSC) are offering such services already; the Virtual Humans predictive modelling platform from ELEMBio allows biotech companies to run supercomputer-based *in silico* trials to assess the cardiac safety ranges of drugs.

A prime example of the use of digital twins in medicine is the virtual human heart (Fig. 1), which is the culmination of research that stretches back more than half a century to experiments in the 1950s on the conveniently-large nerves from a squid. Today, remarkable progress has been made in creating realistic human digital heart. Here we provide but a few examples. Blanca Rodriguez’s team in Oxford makes human virtual heart predictions which are more accurate than comparable animal studies [21], offering a way to reduce vivisection. Our colleagues in the Barcelona Supercomputer Center have developed the Alya Red heart model [22], which typically takes 10 hours to simulate ten heart beats. Working with the company Medtronic, their simulations can help position a pacemaker, fine tune its electrical stimulus and model the effects of an innovative design called the micra [23]. It assesses the

drug dosage and potential interactions between antimalarial drugs to provide guidance for their use in the clinic. Personalised virtual heart models, based on patient data, have been created by a team led by Reza Razavi at King's College London to predict tachycardia [24], while in Johns Hopkins University, a team led by Natalia Trayanova is creating digital replicas of the heart's upper chambers to guide the treatment of irregular heartbeats by the carefully targeted destruction of tissue [25]. In France, Dassault Systèmes has created a cohort of "virtual patients" to help test a synthetic artificial heart valve for regulators, working with the US Food and Drug Administration [26]. This represents another important milestone for Virtual You because, until recently, regulatory agencies have relied on experimental evidence alone.

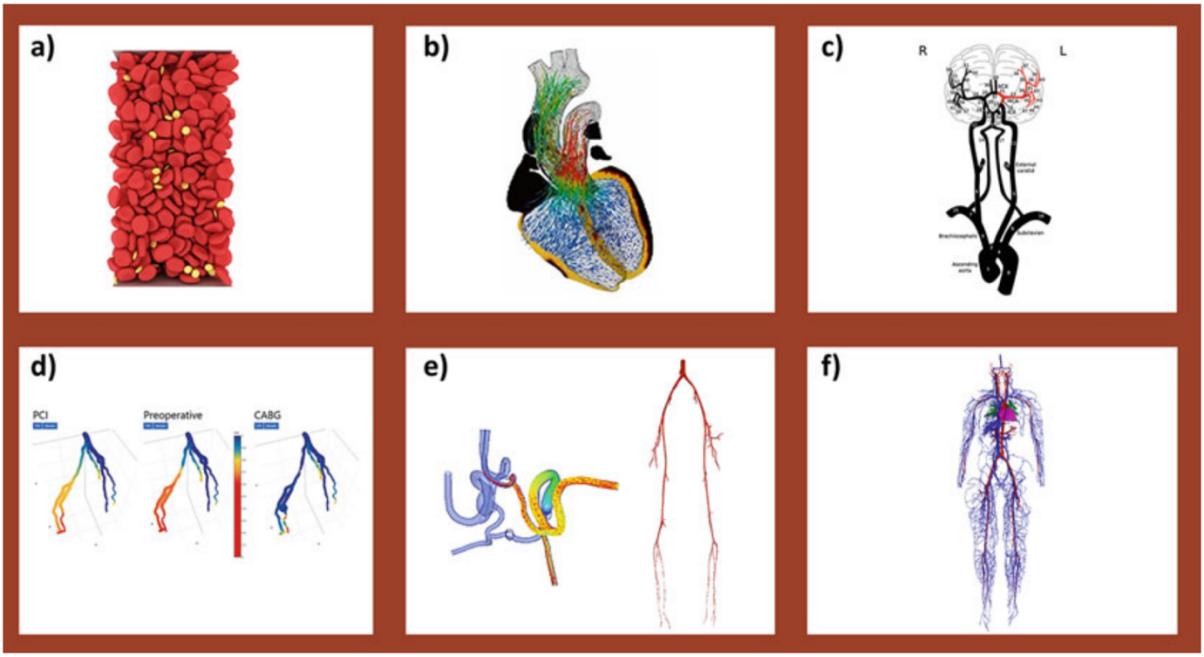

Fig. 1 Digital twin models for the cardiovasculature, which have paved the way for the dawn of precision cardiology. (a) "Digital Blood" and flow-diverting stents within Palabos (University of Amsterdam, Universite´ de Gene` ve, ATOS); (b) Alya Red virtual heart (Barcelona Supercomputing Center, University of Oxford); (c) OpenBF for vascular networks (University of Sheffield, SURFsara); (d) AngioSupport for coronary artery disease (LifeTec Group, SURFsara); (e) HemeLB model of blood flow in arteries (University College London); and (f) HemeLBAlya coupling for cardiovascular flow in virtual human scale geometries (University College London, Barcelona Supercomputing Center).

There are other real and potential clinical applications of digital twins. CT2S is a digital twin workflow that allows prediction of the risk of hip fracture for an individual based on CT scans [27]. The clinical application of this code is to provide a more accurate intervention strategy

for elderly people with weaker bones, such as those who are clinically defined as osteopenic but not receiving any treatments. The workflow provides a complete assessment of bone strength in 3D and analyse the risk factor of that particular individual in sustaining a fall in future. One of us (PVC) has worked with an international team to create a digital twin of a 60,000-mile-long network of vessels, arteries, veins and capillaries using billions of data points from digitized high-resolution cross-sections of a frozen cadaver of a 26-year-old Korean woman, Yoon-sun. By taking over a German supercomputer, SuperMUC-NG for several days, they could show in a realistic way how virtual blood flowed for around 100 seconds through a virtual copy of Yoon-sun's blood vessels down to a fraction of a millimetre across. The team is now charting variations in blood pressure throughout her body and simulating the movement of blood clots. Meanwhile, remarkable progress has been made in creating simple virtual cells [28] by Markus Covert at Stanford University, while at the Auckland Bioengineering Institute Peter Hunter and colleagues are working on an array of organs. Even the most complex known object, the human brain, is being simulated, for instance to plan epilepsy surgery in a French clinical trial. At last, the virtual human is swimming into view.

And there is the remarkable story of how modelling and simulation of the spread of the COVID-19 pandemic was common currency in the UK's national newspapers and media for the best part of two years, being the only means we had available on which to take rational decisions about the kinds of NPI (non-pharmaceutical interventions) government could apply to reduce the spread of covid and the resulting deaths of tens of thousands of citizens. Our own work [29] has revealed the substantial level of uncertainty in one of the most influential models used in Britain (CovidSim).

Quantitative methods and modelling, such as physiologically-based pharmacokinetic (PBPK) modelling, have been involved in the regulatory authorities, which is already well advanced — particular in the FDA. A recent example is the FDA approval of a generic diclofenac sodium topical gel, based on a totality of evidence including a virtual bioequivalence assessment instead of a comparative clinical endpoint bioequivalence study. What matters most here is

how one can define *in silico* protocols that pass muster. Already it is well established that one must conform with specified verification, validation and uncertainty quantification (VVUQ) procedures. V&V has long been required e.g. by the American Institute of Mechanical Engineers, for the design of safety critical devices which engender the safety of human beings. To be used for regulatory decision-making, the models and simulations need to be sufficiently verified and validated (V&V) for their intended purpose.

Summary and perspective

Computational medicine will lead to not only better understanding of mechanisms for disease development and drug treatment but greatly improved healthcare. It should also contribute substantially to reduce the cost of healthcare. It can also accelerate drug development and reduce costs in pharmaceutical industry. Since around 2020, R&D across the pharmaceutical industry has no longer been a profitable activity. The success rate over all drug discovery projects in the industry is no more than about 4% of all those initiated. Computational medicine has potential to change the way drugs are developed in both speeding up the discovery process and reducing the costs, and thus to stimulate drug discovery efforts. Clinically implementing computational models, simulations and digital twins will make medicine truly personalised and predictive.

We still have a long way to go towards developing an AI that comes close to the one between our ears but a good way to start is to place less blind faith in pure data and to give machines more scientific understanding of the world they inhabit. AI, in this ‘smarter’ incarnation that we call ‘Big AI’, features in a new book we have written, Virtual You, the first general account of how to make medicine truly predictive and personal by the use of digital twins - computer models that behave just like an individual person’s body. AI will play a role in this venture, since it has come a long way in the past decade. But we should never forget that we have seen several false dawns during its rise over the past half century (aficionados talk of ‘AI winters’). Today we are enjoying an AI summer, where machines excel over humans in very narrow domains, from playing Go to figuring out the structure of proteins in the body.

But, amid all the hype, we should remember that computers have always been our superior in a narrow way: even the very first such devices could multiply and divide numbers with an ease beyond the ability of most people. Indeed, the current generation of AI is only as good as the data it has been trained on. Give it biased data, and you get biased answers because it makes statistical inferences, glorified curve fitting of these data.

There is also a risk of what is called overfitting – you can think of it as over-learning so while AI can be so in thrall to what it was trained on it can't reliably make sense – or generalise – its predictions to new data and circumstances, such as a new patient it has never encountered before. Even when AI works, after expending a lot of energy and compute power, we have little idea *how* it works: a machine learning algorithm trained to distinguish a chihuahua from a muffin would have no idea what either is.

We need Big AI because, when our virtual twin gives insights into the right drug, lifestyle or diet, we need to understand how it came to its conclusions. In *Virtual You* [20], we discuss various examples of a new generation of 'Big AI', where the brute force of machine learning is augmented with physics based models, mathematical theories which describe how the world actually works. Big AI will play its role in the virtual human and also marks another step towards general artificial intelligence, when an AI agent is finally able to match us in any intellectual task.

Acknowledgments

We are grateful for funding from the UK MRC Medical Bioinformatics project (grant no. MR/L016311/1), the EPSRC funded UK Consortium on Mesoscale Engineering Sciences (UKCOMES grant no. EP/L00030X/1), the European Commission for EU H2020 CompBioMed2 Centre of Excellence (grant no. 823712) and EU H2020 EXDCI-2 project (grant no. 800957). We thank Dr Roger Highfield for his helpful discussions with PVC which led to some of the notions we refer to in this article.

References

1. DiMasi JA, Grabowski HG, Hansen RW (2016) Innovation in the pharmaceutical industry: New estimates of R&D costs. J Health Econ 47:20-33. doi:10.1016/j.jhealeco.2016.01.012

2. Wright DW, Wan S, Shublaq N, Zasada SJ, Coveney PV (2012) From base pair to bedside: molecular simulation and the translation of genomics to personalized medicine. Wiley Interdiscip Rev Syst Biol Med 4 (6):585-598. doi:10.1002/wsbm.1186

3. Wan S, Kumar D, Ilyin V, Al Homsi U, Sher G, Knuth A, Coveney PV (2021) The effect of protein mutations on drug binding suggests ensuing personalised drug selection. Sci Rep 11 (1):13452. doi:10.1038/s41598-021-92785-w

4. McCammon JA, Gelin BR, Karplus M (1977) Dynamics of folded proteins. Nature 267 (5612):585-590. doi:10.1038/267585a0

5. Casalino L, Dommer AC, Gaieb Z, Barros EP, Sztain T, Ahn S-H, Trifan A, Brace A, Bogetti AT, Clyde A, Ma H, Lee H, Turilli M, Khalid S, Chong LT, Simmerling C, Hardy DJ, Maia JDC, Phillips JC, Kurth T, Stern AC, Huang L, McCalpin JD, Tatineni M, Gibbs T, Stone JE, Jha S, Ramanathan A, Amaro RE (2021) AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics. The International Journal of High Performance Computing Applications 35 (5):432-451. doi:10.1177/10943420211006452

6. Libotte GB, Lobato FS, Platt GM, Silva Neto AJ (2020) Determination of an optimal control strategy for vaccine administration in COVID-19 pandemic treatment. Computer Methods and Programs in Biomedicine 196:105664. doi:10.1016/j.cmpb.2020.105664

7. de Ruiter A, Oostenbrink C (2020) Advances in the calculation of binding free energies. Curr Opin Struct Biol 61:207-212. doi:10.1016/j.sbi.2020.01.016

8. Kollman PA, Massova I, Reyes C, Kuhn B, Huo S, Chong L, Lee M, Lee T, Duan Y, Wang W, Donini O, Cieplak P, Srinivasan J, Case DA, Cheatham TE, 3rd (2000) Calculating structures and free energies of complex molecules: combining molecular mechanics and continuum models. Acc Chem Res 33 (12):889-897. doi:10.1021/ar000033j

9. Wan S, Knapp B, Wright DW, Deane CM, Coveney PV (2015) Rapid, Precise, and Reproducible Prediction of Peptide-MHC Binding Affinities from Molecular Dynamics That Correlate Well with Experiment. J Chem Theory Comput 11 (7):3346-3356. doi:10.1021/acs.jctc.5b00179

10. Bhati AP, Wan S, Wright DW, Coveney PV (2017) Rapid, Accurate, Precise, and Reliable Relative Free Energy Prediction Using Ensemble Based Thermodynamic Integration. J Chem Theory Comput 13 (1):210-222. doi:10.1021/acs.jctc.6b00979

11. Bhati AP, Wan S, Alfe D, Clyde AR, Bode M, Tan L, Titov M, Merzky A, Turilli M, Jha S, Highfield RR, Rocchia W, Scafuri N, Succi S, Kranzlmuller D, Mathias G, Wifling D, Donon Y, Di Meglio A, Vallecorsa S, Ma H, Trifan A, Ramanathan A, Brettin T, Partin A, Xia F, Duan X, Stevens R, Coveney PV (2021) Pandemic drugs at pandemic speed: infrastructure for accelerating COVID-19 drug discovery with hybrid machine learning- and physics-based simulations on high-performance computers. Interface Focus 11 (6):20210018. doi:10.1098/rsfs.2021.0018

12. von Delft F, Calmiano M, Chodera J, Griffen E, Lee A, London N, Matviuk T, Perry B, Robinson M, von Delft A (2021) A white-knuckle ride of open COVID drug discovery. Nature 594 (7863):330-332. doi:10.1038/d41586-021-01571-1

13. Spear BB, Heath-Chiozzi M, Huff J (2001) Clinical application of pharmacogenetics. Trends Mol Med 7 (5):201-204. doi:10.1016/s1471-4914(01)01986-4

14. Wan S, Coveney PV (2011) Rapid and accurate ranking of binding affinities of epidermal growth factor receptor sequences with selected lung cancer drugs. J R Soc Interface 8 (61):1114-1127. doi:10.1098/rsif.2010.0609

15. Bunney TD, Wan S, Thiyagarajan N, Sutto L, Williams SV, Ashford P, Koss H, Knowles MA, Gervasio FL, Coveney PV, Katan M (2015) The Effect of Mutations on Drug Sensitivity and Kinase Activity of Fibroblast Growth Factor Receptors: A Combined Experimental and Theoretical Study. EBioMedicine 2 (3):194-204. doi:10.1016/j.ebiom.2015.02.009

16. E. Ellsworth R, J. Decewicz D, D. Shriver C, L. Ellsworth D (2010) Breast Cancer in the Personal Genomics Era. Current Genomics 11 (3):146-161. doi:10.2174/138920210791110951

17. Choudhary A, Fox G, Hey T (eds) (2023) AI for Science. World Scientific Press, Singapore

18. Liu X, Faes L, Kale AU, Wagner SK, Fu DJ, Bruynseels A, Mahendiran T, Moraes G, Shamdas M, Kern C, Ledsam JR, Schmid MK, Balaskas K, Topol EJ, Bachmann LM, Keane PA, Denniston AK (2019) A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digit Health 1 (6):e271-e297. doi:10.1016/S2589-7500(19)30123-2

19. Richens JG, Lee CM, Johri S (2020) Improving the accuracy of medical diagnosis with causal machine learning. Nat Commun 11 (1):3923. doi:10.1038/s41467-020-17419-7
20. Coveney PV, Highfield R (2023) Virtual You: How Building Your Digital Twin Will Revolutionize Medicine and Change Your Life. Princeton University Press,
21. Passini E, Britton OJ, Lu HR, Rohrbacher J, Hermans AN, Gallacher DJ, Greig RJH, Bueno-Orovio A, Rodriguez B (2017) Human In Silico Drug Trials Demonstrate Higher Accuracy than Animal Models in Predicting Clinical Pro-Arrhythmic Cardiotoxicity. Frontiers in Physiology 8
22. Vázquez M, Arís R, Aguado-Sierra J, Houzeaux G, Santiago A, López M, Córdoba P, Rivero M, Cajas JC Alya Red CCM: HPC-Based Cardiac Computational Modelling. In: Klapp J, Ruíz Chavarría G, Medina Ovando A, López Villa A, Sigalotti LDG (eds) Selected Topics of Computational and Experimental Fluid Mechanics, Cham, 2015// 2015. Springer International Publishing, pp 189-207
23. Roberts PR, Clementy N, Al Samadi F, Garweg C, Martinez-Sande JL, Iacopino S, Johansen JB, Vinolas Prat X, Kowal RC, Klug D, Mont L, Steffel J, Li S, Van Osch D, El-Chami MF (2017) A leadless pacemaker in the real-world setting: The Micra Transcatheter Pacing System Post-Approval Registry. Heart Rhythm 14 (9):1375-1379. doi:10.1016/j.hrthm.2017.05.017
24. Chen Z, Cabrera-Lozoya R, Relan J, Sohal M, Shetty A, Karim R, Delingette H, Gill J, Rhode K, Ayache N, Taggart P, Rinaldi CA, Sermesant M, Razavi R (2016) Biophysical Modeling Predicts Ventricular Tachycardia Inducibility and Circuit Morphology: A Combined Clinical Validation and Computer Modeling Approach. Journal of Cardiovascular Electrophysiology 27 (7):851-860. doi:10.1111/jce.12991
25. Boyle PM, Zghaib T, Zahid S, Ali RL, Deng D, Franceschi WH, Hakim JB, Murphy MJ, Prakosa A, Zimmerman SL, Ashikaga H, Marine JE, Kolandaivelu A, Nazarian S, Spragg DD, Calkins H, Trayanova NA (2019) Computationally guided personalized targeted ablation of persistent atrial fibrillation. Nature Biomedical Engineering 3 (11):870-879. doi:10.1038/s41551-019-0437-9
26. Baillargeon B, Rebelo N, Fox DD, Taylor RL, Kuhl E (2014) The Living Heart Project: A robust and integrative simulator for human heart function. European Journal of Mechanics - A/Solids 48:38-47. doi:10.1016/j.euromechsol.2014.04.001
27. Benemerito I, Griffiths W, Allsopp J, Furnass W, Bhattacharya P, Li X, Marzo A, Wood S, Viceconti M, Narracott A (2021) Delivering computationally-intensive digital patient applications to the

clinic: An exemplar solution to predict femoral bone strength from CT data. Computer Methods and Programs in Biomedicine 208:106200. doi:10.1016/j.cmpb.2021.106200

28. Maritan M, Autin L, Karr J, Covert MW, Olson AJ, Goodsell DS (2022) Building Structural Models of a Whole Mycoplasma Cell. J Mol Biol 434 (2):167351. doi:10.1016/j.jmb.2021.167351

29. Edeling W, Arabnejad H, Sinclair R, Suleimenova D, Gopalakrishnan K, Bosak B, Groen D, Mahmood I, Crommelin D, Coveney PV (2021) The impact of uncertainty on predictions of the CovidSim epidemiological code. Nature Computational Science 1 (2):128-135. doi:10.1038/s43588-021-00028-9