Mechanistic models versus machine learning, a fight worth fighting for the biological community?

Ruth Baker *(University of Oxford, Mathematical Institute and St Hugh's College)*

Jose-Maria Pena *(Lurtis Ltd.)*

Jayaratnam Jayamohan *(Oxford University Hospitals, John Radcliffe Hospital, Department of Neurosurgery)*

Antoine Jerusalem *(University of Oxford, Engineering Science and St Hugh's College)*

90% of the world's data have been generated in the last five years [1]. A small fraction of these data is collected with the aim of validating specific hypotheses. These studies are led by the development of mechanistic models focussed on the causality of input-output relationships. However, the vast majority is aimed at supporting statistical or correlation studies that bypass the need for causality and focus exclusively on prediction. Along these lines, there has been a vast increase in the use of machine learning models, in particular in the biomedical and clinical sciences, to try and keep pace with the rate of data generation. Recent successes now beg the question of whether mechanistic models are still relevant in this area. Said otherwise, why should we try to understand the mechanisms of disease progression when we can use machine learning tools to directly predict disease outcome?

Machine learning models provide predictions on the outcomes of complex mechanisms by ploughing through databases of inputs and outputs for a given problem. Authors such as Tom Mitchell define machine learning as the computer-based process that improves its performance on a given task with the experience of it [2]. This experience may come from interactions with previously collected data or from an interaction with the environment while performing the task. Whilst this approach strongly emphasises the artificial intelligence perspective of machine learning, other fields, such as the database community, view machine learning as the algorithmic part of a much broader "knowledge discovery from databases" process (a rather outdated term now known as data mining) [3].

To provide accurate predictions, machine learning models require large amounts of data or an intensive interaction with the environment, the choice of an adequate algorithm, and the identification of inputs and outputs of interest. The ability to avoid the need to understand complex mechanisms, through the use of large scale data sets, engenders machine learning algorithms scalable and efficient in making predictions in e.g. clinical settings. A recent clinical example is the application of Google's DeepMind technology to 1.6 million NHS patient records. This initiative led to the development of the smartphone app Streams,

aimed at addressing “failure to rescue”, where warning signs of deteriorating health are not identified or acted upon quickly enough [4].

While machine learning models can be used to isolate relevant inputs from big data sets for a given output, mechanistic modelling relies on the generation of novel hypotheses for causal mechanisms that are generated through observations of the phenomenon of interest. Its purpose is to mimic real-life events through assumptions on the prominent underlying mechanisms. Typically, this involves constructing simplified mathematical formulations of causal mechanisms, and developing and/or using analytical tools to determine whether the range of possible input-output behaviours predicted by the model, and hence the causal hypotheses, are consistent with experimental observations. Like machine learning models, mechanistic modelling relies upon a two-stage process: first a subset of the available data is used to construct and calibrate the model; and subsequently, in a validation phase, further data is used to confirm and/or refine the model, thereby increasing its accuracy. Ideally the resulting mechanistic models can be leveraged for subsequent use in applications where experiments are either impossible or difficult to achieve.

A paradigm of the mechanistic modelling approach is provided by the work of Hodgkin and Huxley, first published in 1952 [5]. Hodgkin and Huxley’s model of the generation of the nerve action potential is one of the most successful mathematical models of a complex biological process that has ever been formulated. Their model accounts phenomenologically for the dynamics of independent ion channels in which currents are carried entirely by ions moving down electrochemical gradients. It was calibrated and validated using a series of experiments to determine the macroscale parameters of the ion channels (e.g. conductances, equilibrium potentials, the dynamics of each type of ion channel). Hodgkin and Huxley’s work, which won them the Nobel Prize in 1963, has seen widespread use, pushing forward the boundaries of our understanding of the electrical activity of cells on scales ranging from single-celled organisms right through to the neurons in our brains as well as in cardiac mechanics. Hodgkin and Huxley were able to provide scientists with a basic understanding of how nerve cells work. In addition, their model stimulated a significant amount of research in applied mathematics through the derivation of simple caricature models of excitable systems.

Machine learning and mechanistic modelling approaches rely on different types of data and provide access to different types of information, see Table 1. In short, they are two different paradigms. We suggest that, in this sense, they should not be seen as direct competitors or one used at the direct exclusion of another. While mechanistic models provide the causality missing from machine learning approaches, their oversimplified assumptions and extremely specific nature prohibit the universal predictions achievable by machine learning. However, the pros of one are the cons of the other, which suggests that research efforts should be directed towards enabling a symbiotic relationship between both. Returning to the previous example, Hodgkin and Huxley focussed on developing a quantitative model of the action potential produced by the squid giant axon, and multiple variations of this model have been subsequently developed to describe the dynamics of a large number of ion channels.

However, efforts to calibrate Hodgkin-Huxley-type models to all possible neurons remain quixotic as this approach cannot be sustained against the recent acceleration in discoveries of new types of ion channels [6]. This is very much evidenced by the barriers to progress encountered by the existing worldwide initiatives in neuroscience. However, one can imagine in the future the realisation of a universal model applicable to all ion channels where individual parameter calibration and model refinement would rely on a machine learning layer overtop the mechanistic framework. More generally, machine learning research could be harnessed to overcome the current scalability limitations of mechanistic modelling, while mechanistic models could be used by machine learning algorithms both as transient inputs and a validating framework.

However, such marriage is not always straightforward, especially in the clinical context. Any attempt by machine learning technologies to predict individual patient outcomes from past observations using a patient database is potentially able to identify which of existing treatments is most adequate, but intrinsically unable to suggest new treatment protocols or to provide accurate predictions for new treatments. In the literature, this aspect is referred as the “inductive capability” of the learning algorithms (from past data, one can identify patterns happening in the data). This is vastly different from the deductive capability of mechanistic models in which the combination of logical (mechanistic) principles enables extrapolation to predictions about behaviours not present in the original data [7]. In short, mechanistic models can provide insights and understanding into the mechanistic functions of treatments, and these are necessary to overcome the limitations of machine learning predictions. A recent example is the use of machine learning in predicting the success rate for endoscopic third ventriculostomy (used to treat hydrocephalus) [8]. While the algorithm could predict the success rate of the actual procedure, it was not able to take into account the risks for a particular patient of other general physiological variables to allow the procedure to be favoured over another. The exact risks for any patient depend on so many other variables that a clinical estimate of which procedure is preferable can only be made at the bedside. To be able to compare the global state of the patient, the future risks and the treatments necessary requires a “holistic” understanding, and it cannot be provided using machine learning models built on patient data alone.

While the field of fundamental cell biology underpins all advances in our understanding of disease, it remains virtually unaffected by progress in machine learning technologies. One of the reasons for this is that it has traditionally relied on low-throughput means of generating data with a focus on establishing small numbers of input-output relationships. Mechanistic hypotheses generated from the identification of these relationships can often be naturally interrogated using simplified mechanistic models, and causal mechanisms established. More recently, the explosion in high-throughput methods of data collection has reinforced the, often negatively perceived, “butterfly-collecting” nature of cell biology. Yet, the community remains focussed on establishing mechanistic understanding. This dichotomy between data and purpose often leads to the development of a plethora of potential mechanistic models that explain small pieces of a much bigger picture. While these mechanistic models could potentially be assembled as inputs of larger machine learning algorithm, serious efforts in this direction are still lacking.

There are two synergic ways in which mechanistic and machine learning approaches may be combined, see Figure 1. Firstly, within the mechanistic modelling approach, specific components are learnt from the data [9]. An example is the use of multiscale simulations using surrogate models: a machine learning model obtained from the data produced by detailed simulations. Here, only a small number of simulations of the detailed model are run (in order to learn the model), and then the approximate surrogate model is used for future predictions. This approach is used mainly as a method to speed up expensive, computational, multiscale simulations. Secondly, within a machine learning-based pipeline, input information is raw data enriched by derived parameters generated by a mechanistic approach (in the same way improved probabilistic models get the benefit of more informative estimations of the "hidden variables" [10]).

The integration of machine learning approaches and mechanistic modelling in cell biology can be found, for example, in the use of multivariate information measures such as partial information decomposition to identify putative functional relationships between genes from single-transcriptomic data [11]. The validity of the learnt (hypothesised) gene regulatory networks can then be tested using mechanistic modelling approaches, as part of the test-predict-refine-predict cycle so essential in the biological sciences. An example of the use of this approach is in the identification of how pluripotency regulatory networks are reconfigured during the early stages of embryonic stem cell differentiation [12]. However, there remains much work to be done in bringing together machine learning and mechanistic modelling approaches to best effect in the biological sciences.

The lack of progress towards integrating machine learning and mechanistic modelling approaches in fundamental biology studies is at odds with the widespread adoption of machine learning approaches by the clinical community. In fact, many new models and algorithms naturally find their first application in clinical scenarios. This discrepancy of collaboration certainly stems from the relative ease with which one can define the inputs and outputs of a clinical problem, as opposed to the complex fundamental problems tackled by cell biologists. Thus, many of the fast-changing computational approaches spearheaded by interdisciplinary collaborations in the clinical sciences end up escaping entirely this field. Fundamental biology should not choose between small-scale mechanistic understanding and large-scale prediction. It should embrace the complementary strengths of mechanistic modelling and machine learning approaches to provide, for example, the missing link between patient outcome prediction and the mechanistic understanding of disease progression. The training of a new generation of researchers versatile in all these fields will be vital in making this breakthrough. Only then can mechanistic models in cell biology find their real clinical use in the high-throughput world of the 21st century.

Acknowledgment

The authors thank Kate Ellis for her help in designing the figure.

References

1. *Machine learning: the power and promise of computers that learn by example.* Royal Society Report. Issued: April 2017. DES4702 ISBN: 978-1-78252-259-1. Available at https://royalsociety.org/topics-policy/projects/machine-learning/
2. T. M. Mitchell. *Machine Learning.* McGraw-Hill Series in Computer Science. 1997.
3. U. Fayyad, G. Piatetsky-Shapiro and P. Smyth (1996). From data mining to knowledge discovery in databases. *AI Magazine* 17(3)37-54.
4. https://deepmind.com/applied/deepmind-health/working-nhs/how-were-helping-today/
5. A. L. Hodgkin and A. F. Huxley (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology* 117:500-544.
6. http://channelpedia.epfl.ch
7. G. Piatetsky-Shapiro (1996). Advances in knowledge discovery and data mining (Vol. 21). U. M. Fayyad, P. Smyth and R. Uthurusamy (Eds.). Menlo Park: AAAI Press.
8. P. Azimi and H. R. Mohammadi (2014). Predicting endoscopic third ventriculostomy success in childhood hydrocephalus: an artificial neural network analysis. *Journal of Neurosurgical Pediatrics* 13:426-432.
9. C. Xu, C. Wang, F. Ji and X. Yuan (2012). Finite-element neural network-based solving 3-D differential equations in MFL. *IEEE Transactions on Magnetics*, 48(12):4747-4756.
10. J. Binder, D. Koller, S. Russell and K. Kanazawa (1997). Adaptive probabilistic networks with hidden variables. *Machine Learning*, 29(2):213-244.
11. T. E. Chan, M. P. H. Stumpf and A. C. Babtie (2017). Gene regulatory network inference from single-cell data using multivariate information measures. *Cell Systems* 5(3):251-267.
12. P. S. Stumpf, R. C. G. Smith, M. Lenz, A. Schuppert, F.-J. Muller, A. C. Babtie, T. E. Chan, M. P. H. Stumpf, C. P. Please, S. D. Howison, F. Arai and B. D. MacArthur (2017). Stem cell differentiation as a non-Markov stochastic process. *Cell Systems* 5(3):268-282.

Table 1: the advantages and disadvantages of mechanistic modelling and machine learning.

Mechanistic modelling	**Machine learning**
Seeks to establish a mechanistic relationship between inputs and outputs	Seeks to establish statistical relationships and correlations between inputs and outputs
Difficult to accurately incorporate information from multiple space and time scales	Can tackle problems with multiple space and time scales
Capable of handling small datasets	Requires large data sets
Once validated, can be used as a predictive tool where experiments are difficult or costly to perform	Can only make predictions that relate to patterns within the data supplied

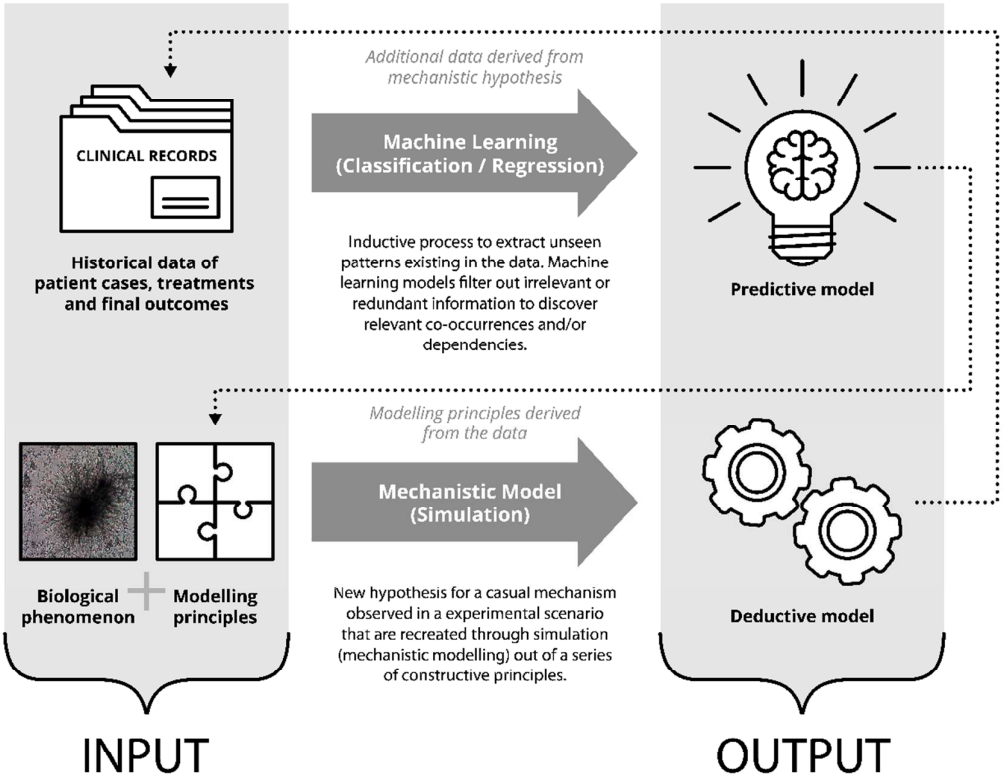

193

194 *Figure 1: The inputs and outputs from machine learning and mechanistic modelling*
195 *approaches, and the potential for synergy between the two.*

196

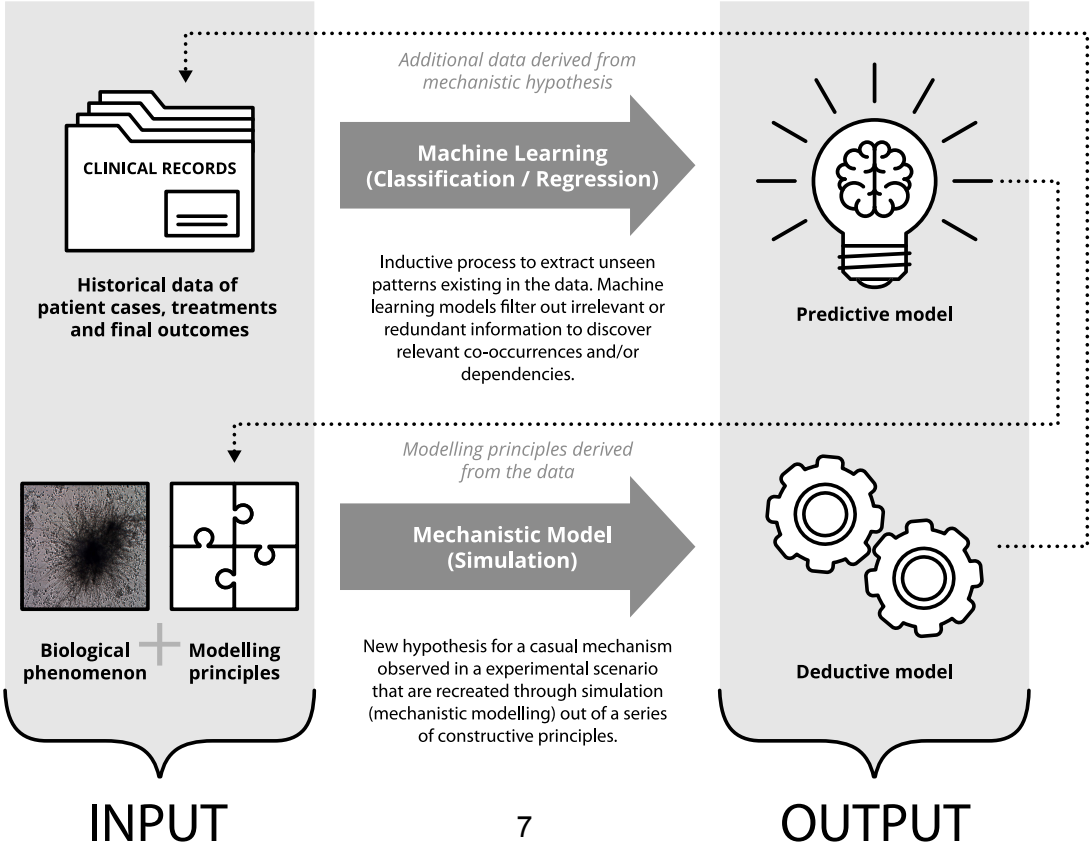
CLINICAL RECORDS
Historical data of patient cases, treatments and final outcomes
Additional data derived from mechanistic hypothesis
Machine Learning (Classification / Regression)
Inductive process to extract unseen patterns existing in the data. Machine learning models filter out irrelevant or redundant information to discover relevant co-occurrences and/or dependencies.
Predictive model
Modelling principles derived from the data
Mechanistic Model (Simulation)
New hypothesis for a casual mechanism observed in a experimental scenario that are recreated through simulation (mechanistic modelling) out of a series of constructive principles.
Deductive model
Biological phenomenon
Modelling principles
INPUT
OUTPUT