THE ROLE OF ARTIFICIAL INTELLIGENCE WITHIN IN SILICO MEDICINE

VPH Institute
Building the Virtual Physiological Human

Avicenna Alliance
Association for Predictive Medicine

Index

Authors

Liesbet Geris, PhD
University of Liège & KU Leuven, Belgium;
VPH Institute; Avicenna Alliance

Cécile F. Rousseau, PhD
Voisin Consulting Life Sciences; Avicenna Alliance

Jerome Noailly, PhD
BCN MedTech, Universitat Pompeu Fabra (UPF), Spain;
VPH Institute; Avicenna Alliance

Payman Afshari, PhD
Johnson and Johnson; Avicenna Alliance

Michaël Auffret, MSc
Voisin Consulting Life Sciences; Avicenna Alliance

Wen-Yang Chu, MSc
Virtonomy.io; Avicenna Alliance

Martha De Cunha-Burgman, MSc
Edwards Lifesciences; Avicenna Alliance

Yoann Lafon, PhD
Université de Lyon, Université Gustave Eiffel, France;
VPH Institute

Thierry Marchal, MSc
Ansys; Avicenna Alliance

Mark Palmer, MD, PhD
Medtronic, plc; Avicenna Alliance

Emmanuelle M. Voisin, PhD
Voisin Consulting Life Sciences; Avicenna Alliance

Marco Viceconti, PhD
Alma Mater Studiorum, University of Bologna,
& IRCCS Istituto Ortopedico Rizzoli, Bologna Italy;
VPH Institute; Avicenna Alliance

Alfons G. Hoekstra, PhD
University of Amsterdam; VPH Institute;
Avicenna Alliance

Gordon Johnston, CEng
Johnson and Johnson; Avicenna Alliance

Rossana Alessandrello, MSc
Agency for Health Quality and Assessment of
Catalonia (AQuAS), Spain; VPH Institute

Antoine Bril, PharmD, PhD
Servier ; Avicenna Alliance

Enrico Dall'Ara, PhD
Insigneo institute for in silico medicine, University of
Sheffield, UK; VPH Institute; Avicenna Alliance

Alexander Kulesza, PhD
Novadiscovery SA; Avicenna Alliance

Raphaëlle Lesage, PhD
VPH Institute; Avicenna Alliance

Bastien Martin, Msc
Novadiscovery SA; Avicenna Alliance

Markus Reiterer, PhD
Medtronic, PLC; Avicenna Alliance

Alicia Waterkeyn
RPP Consulting; Avicenna Alliance

1. Executive summary

The past decade has produced overwhelming evidence that changes in the health status of individuals, measured by well-defined quantitative clinical endpoints, can, in many specific cases, be predicted by computer models (also known as predictive models). This has opened the door to several applications for these computer modelling & simulation technologies, which are generically referred to as *in silico* medicine. There are different ways to build computer models, depending on the technologies as well as the quality and strength of data and knowledge that are used. When models are built using available scientific knowledge in biophysics, biochemistry, and physiology of the human body, both in healthy and diseased states, we refer to them as **knowledge-driven models**. When they are developed directly from data, without making any causal assumption such as is the case for artificial intelligence (AI) methods, we refer to them as **data-driven models**.

AI methods are incredibly powerful, and intense research will make them even better. Although these technologies are very exciting, one should be careful of not over-promising as it could be detrimental to this emerging sector. Knowledge-driven models are neither intrinsically superior nor inherently inferior to data-driven models. What modelling approach is most viable and/or effective in a given situation depends on the question of interest, the context of use and the available data and knowledge. Often, **different *in silico* technologies need to be combined to address very challenging scenarios in healthcare**. After explaining the terminology used in the white paper and the different stakeholders involved, the white paper discusses the overlap and synergies between the different *in silico* medicine approaches (including AI).

Several dimensions must be considered to determine the optimal modelling technology for a defined health scenario. A first dimension is the **availability of reliable mechanistic knowledge** of the phenomenon of interest. In the absence of such mechanistic knowledge, data-driven models are the clearly preferable. A second dimension is the effort, defined as the **computational cost**, meaning conducting the simulation within the time limit imposed by the context of use. The computer models providing a real-time answer require different technologies, computing facilities and deployment strategies than computer models used in the regulatory approval process or for planning medical interventions. A third dimension is linked to the requirements that each modelling strategy has in terms of **quality and quantity of data** required to build, run, and validate each predictive model. The availability of large high-quality data sets remains today an important challenge in healthcare. Additionally, the cost for acquisition, preparation, and management of data will become increasingly significant, particularly for data-driven models. A fourth and last dimension is related to the process we use to **establish the credibility of a prediction** obtained by such models. Both knowledge-driven and data-driven models must undergo complex scrutiny before they can be used in clinical or regulatory practice. However, the extent and the nature of such scrutiny is different for these two types of models as are the risks related to use outside the validation domain and concept drift.

The white paper describes how the greatest public health gain should be obtained by combining various approaches and gives several examples to illustrate the points made. Data-driven elements can be introduced in knowledge-driven models **to complement, assess or accelerate** the models and their computation by simulating parts of the modelled phenomenon for which insufficient mechanistic knowledge exists. When **personalising knowledge-driven models** with patient specific data, the knowledge-driven model needs to be inverted to find the appropriate parameter values and their uncertainties. In many cases, simple fitting techniques will not work properly, and developers will need to resort to advanced data-driven methods. Finally, to meet requirements on accuracy and speed of simulations, the knowledge-driven model can be replaced (completely or part of it), by a data-driven **surrogate model**. The combination of knowledge-driven and data-driven models allows the modelling as a whole to be computationally affordable while maximising the benefits of the knowledge-driven baseline.

In the other direction, mechanistic elements can substantially **augment data-driven models**. To develop data-driven models with strong predictive power, large amounts of data (*i.e.* big data) are necessary. In many medical applications, however, the amount of clinical data available is by far not enough to sufficiently train a data-driven model. A valid knowledge-driven model can be used as **a source of data** to enhance the predictive power of data-driven models, among others for rare events. Alternatively, the explicit inclusion of the knowledge-driven model as prior knowledge will help to **reduce the size of the training data set** required to train the data-driven models. Another way in which knowledge-driven models can be used to enhance data-driven models is to **provide a benchmark** against which the accuracy of the developed algorithms can be tested.

Health policies can lead to the development of powerful *in silico* technologies (in the broadest meaning). The explicit inclusion of modelling in regulatory legislation such as the medical device regulation has given an impetus to the relevant stakeholders to consider modelling and simulation in their respective processes. The European Commission's creation of the European Health Data Space and the proposed AI legislation will have an important impact on the further development of not only data-driven models but of *in silico* medicine as a whole. Reciprocally, *in silico* models, be they knowledge-driven, data-driven or hybrid, can drive health policies.

To date, EU health policy has not put forward any formal regulatory frameworks for the various applications of *in silico* medicine and as such it remains a grey area when it comes to legislative guidance. It is crucial to strive for a **policy framework** that can enable (amongst others) *in silico* medicine in a manner that can raise safety standards in clinical trials, improve the lives of rare disease patients and reduce costs to device manufacturers and researchers. Stakeholders require legal certainty from regulators on what constitutes an acceptable model and as such need a framework which they can rely on to employ these techniques for safer and more effective healthcare delivery.

In line with this context, the members of the Avicenna Alliance from medtech, pharmaceutical, software and life science industries, as well as from the *in silico* medicine academic community and healthcare organisations **call on the European Commission to**:

- Reflect the **added value that *in silico* technologies have in healthcare** within the revision of existing pharmaceutical legislation. This is particularly the case for the production of digital evidence for treatment development, especially for paediatric and rare diseases, which cannot be generated by the *in vitro* or *ex vivo* models, or is very challenging to generate using *in vivo* (animal and patient) models.
- Introduce a clear regulatory framework for digital evidence for all medical products, including both drugs and devices, that takes into account the context of use and model risk to **reduce political and economic barriers.**
- Develop a **Good Simulation Practice**, similar to Good Clinical Practice or Good Manufacturing Practice, as a starting point for developing *in silico* models used in decision-making processes such as regulation.
- Launch **funding schemes** promoting the adoption of *in silico* technologies in healthcare through pre-commercial procurements, public procurements of innovations and early dialogues either with local/national health technology assessment agencies or through EU-level joint clinical assessments.
- **Promote public trust** in *in silico* technologies by increasing awareness, e.g. increasing public literacy about said technologies.

2. Aims and scope

Currently, Artificial Intelligence (AI) and data are receiving significant attention in terms of policy activities and research funding at the European level and across the globe. At the same time, computer modelling and simulation are increasingly mentioned as a separate technology next to artificial intelligence[1,2]. However, there is no clarity as to differences and similarities of these technologies. The Avicenna Alliance (AA) is an association of industry, academia and healthcare organisations, the latter two represented through the Virtual Physiological Human institute (VPHi), that has a strong interest in the development and adoption of *in silico* technologies in all areas of healthcare. **With this white paper, the authors aim to provide clarity on the overlap and synergies between the different *in silico* medicine approaches (including AI), as well as their respective place within the overarching continuum of *in silico* medicine and digital healthcare.**

After an introduction into the different concepts, this paper presents a dedicated *in silico* strategy from the perspective of the users of the technologies. Subsequently, it discusses the *in silico* spectrum from knowledge-driven to data-driven methodologies. Afterwards, an overview is provided of how these *in silico* technologies can be used, in which context, and at what point in time. It then explains how public health can benefit from the combination of different *in silico* approaches rather than from any a priori focus on a single technology, illustrated by several examples. Finally, a number of elements related to the assessment of credibility of the *in silico* evidence are discussed, in the context of health and medicine. The paper ends with a conclusion and recommendations to policy makers, regulators and payers in Europe and beyond.

> ***"This white paper provides clarity on the overlap and synergies between the different in silico medicine approaches, including AI.***

1: Amendments adopted by the European Parliament on 14 September 2021 on the proposal for a regulation of the European Parliament and of the Council amending Regulation (EC) No 851/2004 establishing a European Centre for disease prevention and control (COM(2020)0726 — C9-0366/2020 — 2020/0320(COD)). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021AP0376 .

2: Horizon Europe work programme 2021-2022 for Health. https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/wp-call/2021-2022/wp-4-health_horizon-2021-2022_en.pdf

3. Terminology

This section is not intended to provide an exhaustive lexicon on all terms related to *in silico* technologies, it merely serves to clarify the definition of a number of key concepts used throughout this paper that are not explicitly defined in the body of the paper. In several cases, the definition of a term can differ based on the community that is using it or the context in which it is used. In those cases, several definitions have been provided along with their source.

In silico[3]: *in silico* means carried out in the computer, which is in contrast to *in vitro* (on the bench), *ex vivo* (outside the living organism), or *in vivo* (inside the living organism),

- *In silico* methods are experiments or analyses carried out *in silico*.
- *In silico* medicine encompasses the use of *in silico* technologies in all aspects of the prevention, diagnosis, follow-up, prognostic assessment, and treatment of diseases. These can be generic or individualized models.
- *In silico* trial means the use of individualized computer simulations in the development or regulatory evaluation of a medicinal product, medical device, or medical intervention. An *in silico* trial can apply to nonclinical or clinical studies.
- *In silico* test/experiment/studies means the simulation of an *in vitro*, *ex vivo* or *in vivo*, but not clinical, experiment in the computational environment.
- *In silico* evidence (sometimes also refer to as digital evidence) is the result from *in silico* methods applied in any activity involving a decision on quality (such as a regulatory process), and encompasses *in silico* trials and *in silico* tests.

Computational model, model, simulation:

- Computational modelling is the use of computers to simulate and study real-world systems using mathematics, physics and computer science[4].
- Computer models are the algorithms and equations used on a computer to capture the behaviour of a physical system.
- A computer simulation is the result of running a computer program that contains equations or algorithms. Simulation, therefore, is the process of running a model.

Digital Twin: definitions vary among the use of digital twins in engineering or healthcare contexts.

- In engineering, a digital twin is defined as a set of virtual information constructs that mimic the structure, context, and behaviour of an individual/unique physical asset, process or entity, is dynamically updated with data from its physical twin throughout the period where the twin is used to support decisions, and ultimately informs decisions that realize value.
- In healthcare the term is often used in the context of personalised medicine, as the direct use of individual-specific models for the prevention, prediction, screening, diagnosis and treatment of a disease, as well as the evaluation, optimization, selection and personalisation of intervention options[5]. To acknowledge this is a relaxed version of the original definition of a digital twin, it will be called a digital patient, a virtual patient, a virtual twin, a digital avatar, a human digital twin or a digital twin for personalised medicine.

Knowledge-driven models (used as synonyms: mechanistic models, hypothesis-based models, physics-based models or white box models): *in silico* models that are based on prior knowledge on cause-and-effect relationships[6,7].

Data-driven models (used as synonyms: phenomenological models, empirical models, black box models): *in silico* models that develop a predictor automatically for the data without making any causal assumptions are

---

3: Viceconti M. et al.. *In silico* Clinical Trials: how computer simulations will transform the biomedical industry. 2016, https://avicenna-alliance.com/files/user_upload/PDF/Avicenna_Roadmap.pdf

4: National Institute of Biomedical Imaging and Bioengineering (NIBIB), 2020

5: Consensus definition from the European Commission workshop on the human digital twin December 2020.

6: Transtrum M.K. & Qiu P. (2016). Bridging Mechanistic and Phenomenological Models of Complex Biological Systems. PLoS Comput Biol. 12(5):e1004915. doi: 10.1371/journal.pcbi.1004915.

7: Lema-Perez L. (2019). On parameter interpretability of phenomenological-based semiphysical models in biology. Inform. Med. Unlocked 15, 100158. doi:ff10.1016/j.imu.2019.02.002f

called data-driven models. In this context, models utilizing statistics, artificial intelligence, machine learning, and deep learning are considered data-driven models.

Machine Learning (ML): ML is the scientific study of computer algorithms that are able to learn and adapt through experience. ML algorithms build a model based on training data, in order to make predictions or decisions without being explicitly programmed to do so[8]. There are three main categories of ML algorithms: supervised, unsupervised and reinforced. Within (and across) these categories, different ML models have been developed such as artificial neural networks (which also make up the backbone of the ML subgroup called deep learning[9]), evolutionary algorithms, etc.

Artificial Intelligence (AI): There is a large variety in definitions of AI. We refer the reader to the 2021 Joint Research Centre (JRC) report[8] providing a full taxonomy of AI. Here we cite several definitions currently used by the European Commission and other international organisations.. The common element is the ability to learn from data and adapt its behaviour accordingly.

- European Union (EU) high-level expert group on AI[10]: AI refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals. In the context of this paper, only the purely software-based AI systems are considered (and more specifically, the analytical AI that is focused on cognitive intelligence and decision-making as opposed to human-inspired or humanised AI). Hardware-based AI systems that are also included in this definition are beyond the scope of this paper.
- AI Act (European Commission (EC), 2021)[11]: AI system means software that is developed with one or more of the techniques and approaches mentioned hereafter and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with. The aforementioned techniques and approaches are:
    - **a.** Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
    - **b.** Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;
    - **c.** Statistical approaches, Bayesian estimation, search and optimization methods.
- Organisation for Economic Cooperation and Development (OECD)[12]. An AI system is a machine-based system that is capable of influencing the environment by producing an output (predictions, recommendations or decisions) for a given set of objectives. It uses machine and/or human-based data and inputs to:
    - **a.** Perceive real and/or virtual environments;
    - **b.** Abstract these perceptions into models through analysis in an automated manner (*e.g.,* with machine learning), or manually; and
    - **c.** Use model inference to formulate options for outcomes.

  AI systems are designed to operate with varying levels of autonomy.

Context of Use (CoU): The context of use defines the specific role and scope of the computational model used to address the question of interest[13]. It should include a detailed statement of what will be modelled and how the outputs from the computational model will be used to answer or inform the question of interest. It is important to note that the CoU is distinct from the "indications for use" or "intended use" of a medical device, which are descriptions of how a device is intended to be used in clinical practice.

8: Samoili, S., et al., AI Watch. Defining Artificial Intelligence 2.0, EUR 30873 EN, Publications Office of the European Union, Luxembourg, 2021, ISBN 978-92-76-42648-6, doi:10.2760/019901, JRC126426.

9: Ching, T., D. S. et al. (2018). "Opportunities and obstacles for deep learning in biology and medicine." J R Soc Interface 15(141). doi: 10.1098/rsif.2017.0387

10: https://digital-strategy.ec.europa.eu/en/library/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines

11: European Commission, Communication 2021/0106 (COD): Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts.

12: OECD Framework for the Classification of AI systems, OECD Digital Economy Papers, No. 323, OECD Publishing, Paris, https://doi.org/10.1787/cb6d9eca-en.

13: Assessing Credibility of Computational Modeling through Verification and Validation: Application to Medical Devices V&V 40 – 2018. ASME, 2018. 60p. ISBN: 9780791872048.

Model Credibility: Model credibility refers to the trust in the predictive capability of a computational model for the CoU[12]. Trust can be established through the collection of evidence from the credibility activities. The process of establishing trust includes performing verification, validation and uncertainty quantification (VVUQ) and then demonstrating the applicability of the verification & validation (V&V) evidence to support the use of the computational model for the CoU. This is described in detail in the ASME V&V40 standard[12].

Verification *(in the context of in silico technologies):* Verification is the exercise to ensure that the mathematical model is implemented correctly and then accurately solved[14]. Verification is composed of two activities: code verification and calculation verification. It often encompasses comparison with a reference "source of truth ", like an analytical solution or a convergence limit.

Validation *(in the context of computer modelling and simulation):* Validation is the process of assessing the degree to which the computational model is an appropriate representation of the reality of interest[13] in a specific CoU, assessed using reference data as comparator. Validation is generally demonstrated by comparing the computational model predictions with the results from the comparator(s), which might be *in vitro*, *ex vivo* or *in vivo* data. Therefore, appropriate validation activities require attention to both the computational model and the comparator(s). Differences exist in the exact validation strategies for knowledge-driven vs data-driven models. For instance, a dataset for evaluating of predictive capability of a finalized model would be called a 'validation dataset' in V&V 40 but is called a 'testing dataset' in ML.

Uncertainty Quantification: Uncertainty Quantification (UQ) is the process used to determine how dispersion or lack of exact knowledge of inputs, parameters, processes as well as unaccounted for or random factors affect the output of a model or algorithm. Quantification of uncertainty is essential for the end users to determine the degree of confidence in decisions based on the output of models and algorithms. Therefore the Uncertainty Quantification is part of the Validation process[12].

> ***"Knowledge-driven models are models that are based on prior knowledge on cause-and-effect relationships. [...] Data-driven models are models developed on data without making any causal assumptions. [...] These are the extremes of the in silico spectrum, with most models situated somewhere in between.***

14: Standard for Verification and Validation in Computational Solid Mechanics V&V 10 – 2019. ASME, 2020. 44p. ISBN: 9780791873168.

4. Introduction

Artificial Intelligence is omnipresent in policy and regulation nowadays. At the same time, there is an increasing awareness that other digital technologies can play an important role in healthcare. The last ten years produced overwhelming evidence that, despite their complexity, changes in the health status of individuals as measured by well-defined quantitative clinical endpoints can, in many specific instances, be reproduced and predicted by computer models. This has opened the door to a number of relevant applications for these computer modelling & simulation technologies in healthcare, which are generically referred to as *in silico* medicine.

In silico models are used in software for medical purposes (*e.g.*, decision support systems) as well as in the development and testing of medicinal products and medical devices. The European Parliament Research Service's recent publication via the Scientific Foresight Unit (STOA)[15] recognises that *in silico* testing, along with *in vitro* testing, has great potential in lowering cost of drug discovery; the same study also speaks of *in silico* validation in discrete testing of AI algorithms. Various relevant definitions have been provided in the Terminology section. The exact nature of the computer modelling and simulation depends on the users of the technologies or the evidence they generate.

Health Technology Assessment (HTA) agencies and payers:

HTA is an evidence-based process that allows competent authorities to determine the relative effectiveness of new or existing technologies (medicinal products and medical devices). It focuses specifically on the added value of a health technology in comparison with other new or existing health technologies. Such added value could be assessed using computer modelling and simulation, balancing benefits and harms of the new health technologies in different scenarios or for different patients' profiles, as well as balancing benefits and resources required to achieve them and relevant sources of uncertainty in decision making[16]. Through HTA, national health authorities and payers can take informed decisions on pricing or reimbursement.

Regulatory agencies:

Use of *in silico* evidence generated by models that can be included in regulatory dossiers to demonstrate patient safety and efficacy/performance. Various agencies have specific guidances for the use of computer modelling and simulation in the development of medicinal products from design over first-in-human dosing justification to late-stage confirmatory trial population selection eligibility criteria formulation (*e.g.*, EMA: Model Drug Discovery Development; USA-FDA Model Informed Drug Development).

Healthcare professionals:

Digital Patient solutions (also known as Digital Twins or Digital Avatars) provide valuable support to the medical decision on individual patients, whether related to diagnosis, prognosis or treatment planning, and they enable personalised medicine[17].

Patient:

Personal Health Forecasting solutions supplement the digital health revolution with the vital element of forecasting, as a guide for self-management of chronic patients and people at risk of developing diseases[15].

R&D (Academia, Industry):

In silico Trials use Digital Patients in the place of or in addition to real patients in the testing of safety or efficacy of new medical products (*e.g.,* medicinal product, devices). Their use for healthcare products brings the analogue to virtual testing that is now common practice in other industrial sectors. Without the same limitations as a physical trial, they enable testing of new medical products on a relevant virtual patient population.

15: Artificial intelligence in healthcare: Applications, risks, and ethical and societal impacts, Study, 01-06-2022. ISBN: 978-92-846-9456-3

16: Dahabreh IJ, et al., Modeling and Simulation in the Context of Health Technology Assessment: Review of Existing Guidance, Future Research Needs, and Validity Assessment [Internet]. Rockville (MD): Agency for Healthcare Research and Quality (US); 2017 Jan. Report No.: 16(17)-EHC020-EF. PMID: 28182366.

17: Díaz V., et al., Discipulus Roadmap for the Digital Patient, 2013
https://www.vph-institute.org/upload/discipulus-digital-patient-research-roadmap_5270f44c03856.pdf

5. The *in silico* continuum

To build knowledge in science, we observe natural phenomena and formulate a large number of alternative hypotheses that can explain those observations in terms of cause-effect relationship. Through controlled experiments, we can disprove the vast majority of these hypotheses; however, a few resist any attempt to disprove them, and are thus accepted by the scientific community as tentatively true, and become part of the scientific knowledge (*e.g.*, like the Newton's laws of Dynamics). Hereinafter we will refer to predictive models that are based on such cause-effect relationships with the term **knowledge-driven models**, because they are built on the knowledge of the causal mechanism of the process being modelled. The predictive *in silico* models that are described and tested in the growing scientific literature or that have already reached the market and are in widespread use, are largely based on the available scientific knowledge on biophysics, biochemistry and physiology of the human body both in healthy and diseased states.

Nevertheless, this is not the only way to build predictive models. When a sufficiently large body of empirical observations are available, we can try to develop a predictor without making any causal hypothesis. This can be done in a variety of ways: using statistics, system identification methods, machine learning, etc. Hereinafter we will refer to these other predictive models with the term **data-driven models** because they are derived exclusively by the observation of the phenomenon to be modelled.

It can be convenient to describe complex knowledge spaces with well-defined taxonomies, like the separation between knowledge-driven and data-driven models we use in this paper. But it is important to keep in mind the limitations of such convenient separation. The most important is that purely knowledge-driven models rarely exist; real-world knowledge-driven models always include some data-driven elements. For example, if a model describes the laws of physics for a biological system on a finite portion of space-time, the boundary conditions that describe the effects of the rest of the universe (not explicitly incorporated in the model) on that portion of space-time must necessarily be described with data-driven elements. In addition, when a knowledge-driven model describes events taking place at the level of a tissue or organ, it makes abstraction of the components at the cellular scale and events happening at this cellular scale with impact on the tissue/organ scale are described using data-driven descriptions.

In general, we can say that a predictive model of a biological system is always built in a data-driven way from empirical observations of said biological system, with variable amounts of prior mechanistic knowledge. We call data-driven models those that contain no mechanistic knowledge, and knowledge-driven models those that are largely made of mechanistic knowledge. However, these are the extremes of the ***in silico* spectrum**, with most models situated somewhere in between these extremes as depicted in **Figure 1**.

Lately, the boundaries between these disciplines are blurring with explainable AI[18], mechanistic ML[19] or ML-based surrogates of knowledge-driven models emerging as techniques that combine the best of both worlds - auditable and rationalizable causality with large flexibility and capability to be trained for a large number of applications and outputs.

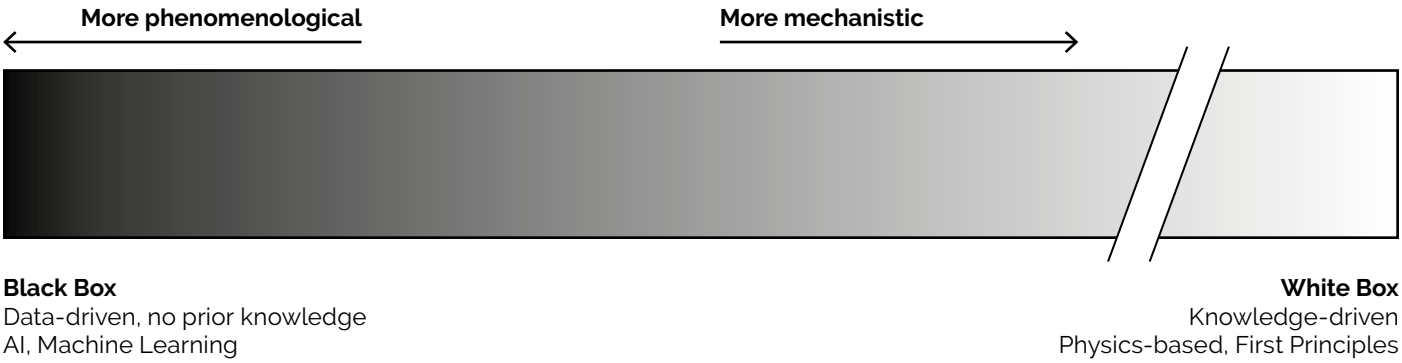

Figure 1: *The in silico spectrum including the most commonly used synonyms for the extremes of the spectrum.*

18: Barredo Arrieta A, et al., Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 2020 ; 58(C) :82–115202. https://doi.org/10.1016/j.inffus.2019.12.012.

19: Hurault G, et al., Personalized prediction of daily eczema severity scores using a mechanistic machine learning model. Clin Exp Allergy. 2020 Nov;50(11):1258-1266. doi: 10.1111/cea.13717.

6. Predictive models in healthcare: what to use and when?

AI methods are amazingly powerful, and there is intense research focusing on making them even better. But, like many other innovative approaches that have appeared in recent years, there is also a tremendous hype for the possibilities that these methods can offer in healthcare, mostly driven by the desire to minimize development times and arrive first at the forefront. One should be careful of not over-promising as it could be detrimental to this emerging sector.

This paper wants to explain the importance of AI methods within in silico medicine, and to recognise the strengths and limitations of knowledge-driven and data-driven models in order to guide their rational and conscious adoption. We also want to highlight the huge potential of the combination of these approaches in addressing very challenging scenarios in healthcare.

Let us start by making a general statement: knowledge-driven models are not inherently superior or inherently inferior to data-driven models. What modelling approach is most viable and/or effective in a given situation depends on the question of interest, the CoU and the available data and knowledge.

A first important dimension to decide what modelling technology to use, is the **availability of reliable knowledge on the phenomenon of interest** and if it is sufficient to formalize a mathematical description. If there is very little, if any, reliable causal knowledge around the biological process we want to model, data-driven models are clearly preferable.

A second dimension of the problem is the effort, defined as the **computational cost**, meaning the cost to obtain the prediction in the time limit imposed by the CoU. A predictive model used to support a medical decision in an emergency room must provide an answer in minutes, or else that prediction is useless; a prediction to be used to plan next week's surgery can be provided in days without any problem. The computer models providing a real-time answer require different technologies, computing facilities and deployment strategies than computer models used in the regulatory approval process or for planning medical interventions. In most cases, we can say that while data-driven models have a very high computational cost when they are being built, they usually have a very small computational cost when they execute. Knowledge-driven models are generally the opposite: their construction does usually involve moderate computational costs but their execution may require medium to high computational costs, especially when the models include detailed mechanistic or personalized descriptions. Surrogate modelling, *i.e.* training fast-running models on output data obtained from extensive knowledge-based model simulations, allows to investigate at low cost the outcomes of such knowledge-driven models, making it suitable for use as medical decision support in an emergency room.

A third dimension is related to the **data requirements** in terms of **quality and quantity of data required** to build, validate, and run each predictive model. Observational data are required to develop a predictive model while other data are required to quantify its predictive accuracy (validation). In the case of ML models, yet another data set is required to test the model after validation of its parameters. However, in order to be specific, either to a single patient, or to small sub-groups of patients with common characteristics, predictive models need also to be fed at run-time with patient-specific data. These data sets (to build, validate and test/run a model) are usually different, both in nature and requirements, for different modelling technologies. Sometimes we have appropriate data to build, validate or run knowledge-driven models but not data-driven ones, or vice versa. One of the biggest challenges in the development and use of *in silico* models in healthcare is the **availability of and access to data**. The absence of an abundance of clean and consistent data is a challenge setting it aside from the other applications where *in silico* models are used in industry and society. Healthcare data is not free, both in terms of the cost associated with the acquisition of it and in terms of access due to privacy restrictions that apply (*e.g.,* General Data Protection Regulation (GDPR) restrictions related to patient consent in EU). Additionally, the clean-up of the data to collect, store, get it into a usable format and manage, it comes with a non-negligible cost. Although data requirements are not the same for knowledge-driven and data-driven models, both require the presence of reliable data to establish the model and assess its credibility. For knowledge-driven models, the quality of the mechanisms/knowledge used to build the model is also dependent on the quality of the data this knowledge was derived from.

4

A fourth dimension is related to the process we use to **establish the credibility of a prediction obtained by such models**. Both knowledge-driven and data-driven models need to undergo complex scrutiny before they can be used in the clinical or regulatory practice, but the extent and the nature of such scrutiny is different for the two types. Two topics are important to mention in this context: predictions outside of the validation domain and concept drift. The former refers to the use of a model for situations that fall outside of the domain for which the model has been validated (*e.g.,* application of models validated for adult cases to paediatric cases). The latter indicates that the relationship between input and output variables can change over time in unforeseen ways (*e.g.,* due to disease progression or intervention variability) potentially causing the model to become less and less effective over time (*i.e.* running the simulations for longer than the duration for which validation data was available).

These are risks that could be said to exist for both data-driven and knowledge-driven models. However, knowledge-driven models are built on causational knowledge that generally reflects robust foundational mechanisms and tend to be relatively invariant when **operating outside of the validation domain** or when disease progresses, therefore we can, with relative confidence, confirm the credibility of knowledge-driven models "once for all" in the validation phase providing there is sufficient rigor in the verification and validation activities, outcomes and use. Data-driven models, on the other hand, are more susceptible to these recognized risks and in some cases such extrapolation might not even be possible. These models are often wholly correlational and founded on how "representative" data features and data observations used to construct the model are relative to current and future patterns of data variation in the intended use application. For data-driven models, mitigating the risk of predicting beyond the domain of validation is a question of clearly defining that validation space when aligning on the CoU at the very beginning, applying appropriate rigor in feature selection and representative data collection (pre-validation), and sufficiently communicating that CoU and the limitations on prediction applicability throughout the model lifecycle to all stakeholders (post-validation).

To mitigate the risk of **concept drift** with data driven models, similar to what was mentioned for the validation domain, this involves a conversation on post-validation topics such as meaningful performance monitoring and robust change control mechanisms. Performance monitoring of the model is relatively self-explanatory and needs the engagement of both provider and user. Change control is a conversation on either "locking" models and changing them at some frequency, or as a result of some performance trigger, or allowing the models to continuously adapt to new unseen data. The change control approach chosen should at least be informed by benefit-risk and the ease of model adaptation. Care has to be taken with both scenarios. Locked models can experience concept drift relatively quickly if the change frequency is too long and models that adapt continuously are open to rare event data. Given the inherent risks with correlational analytics, developers of data-driven models should always be open to challenging ground truth beliefs and the potential for new data that challenges the interpolation or extrapolation capability of the model. Along with those post-validation controls, data collection and data analysis should also be recognized as lifecycle activities with such models.

> ***"What modelling approach is most viable and/or effective in a given situation depends on the question of interest, the context of use as well as the available data and knowledge.***

7. Combining knowledge-driven and data-driven approaches

Most of the predictive models that are developed in healthcare are not situated at the extremes of the *in silico* spectrum, they rather are a combination of knowledge-driven and data-driven elements. In fact, when used in concert both approaches can strengthen each other, to the extent that they can push *in silico* medicine applications towards true clinical practice. Below details are provided on the different ways knowledge-driven and data-driven approaches can be combined. Three tangible examples are summarized in pages 16 to 18 to help illustrate the various points made.

There are multiple ways in which **data-driven elements can be introduced in knowledge-driven models**, other than through *boundary conditions* mentioned above. For instance, every knowledge-driven model is equipped with parameters (*e.g.*, mechanical properties of a tissue that is modelled), but these parameters are typically quite variable between individuals. When *personalising knowledge-driven models with patient specific data*, we need to find the appropriate parameter values and their uncertainties. In many cases, simple fitting techniques will not work properly due to the limited amount of data available, and we need to resort to advanced data-driven methods (such as evolutionary algorithms or Bayesian inference). A good example of this is the personalisation of models describing drug absorption, distribution, metabolism and excretion (so-called physiology-based pharmacokinetic models) using clinical data[20]. Data-driven modelling can also be used to efficiently explore the *variability and error propagation due to parameter uncertainties* in knowledge-driven models, especially when those are applied to biological materials and processes. For instance, when modelling a stent, the model of the stent itself has less uncertainties than the model of the aortic wall the stent is in contact with. Surrogate modelling techniques are typically used to *explore the most relevant variations of model prediction* and target the most critical data-driven parameters that can be used as primary calibration or design variables, or point out the need for specific experimental measurements. Finally, data-driven modelling can be used to *identify specific subgroups in clinical data*, after which knowledge-driven models can be used to identify the root cause mechanisms in each of the different subgroups (**example 1**, p. 17).

In knowledge-driven models there is always a trade-off between the accuracy of the prediction and the required computational time. Generally, an increase in accuracy will lead to an increase in computational time. If, however, the predication needs to be delivered in a pre-set time, which in acute situations or during a surgery can be quite short, it might be that a knowledge-driven model cannot deliver an accurate enough prediction. One tried and tested solution is to *replace the knowledge-driven model, or part of it, by a surrogate model*. The basic principle of surrogate models is that these are data-driven models that are trained using data obtained from simulations executed with the knowledge-driven model (see **example 2**, p. 18). By doing this, the surrogate model can be trained 'off-line', and then applied in the pre-set time scenarios. Many knowledge-driven models combine sub-models for different parts of the physiology (*e.g.,* one sub-model for blood flow, another for tissue growth or thrombus formation[21]). Typically, in such combined models, a major part of the computing time is taken up by one or two of the sub-models. Replacing these sub-models by surrogates results in hybrid solutions, mixing knowledge-driven and data-driven models, that would render the whole affordable computationally while still maximally benefiting from the advantages of the knowledge-driven baseline.

In the other direction, **knowledge-driven elements can substantially augment data-driven models**. To develop data-driven models with strong predictive power, one needs sufficient high-quality data. In many medical applications, however, the amount of available high-quality clinical data is not enough to sufficiently train a data-driven model and does not allow for new machine learning techniques to augment the data set. For example, in the case of image analysis applications, training can be performed predominantly on patients with a disease because these are the most common scans to be performed. The result can be that the data-driven model always classifies images as diseased even if an image is taken from a healthy person. Alternatively, a classifier may not work on edge cases, this could be images from children, very large or very small people, or people with congenital diseases, co-morbidities, surgeries, tumours, or implanted devices. The myriad of possibilities poses very real limits on developing general clinical data-driven applications in cases where vast general databases are not available. Knowledge-based model simulations can be used very effectively to provide

20: Tsamandouras N, et al., Combining the 'bottom up' and 'top down' approaches in pharmacokinetic modelling: fitting PBPK models to observed clinical data. Br J Clin Pharmacol. 2015;79(1):48-55. doi: 10.1111/bcp.12234.

21: Nikishova A, et al., Semi-intrusive multiscale metamodelling uncertainty quantification with application to a model of in-stent restenosis. Philos Trans A Math Phys Eng Sci. 2019;377(2142):20180154. doi: 10.1098/rsta.2018.0154.

richer data sets of edge cases, to understand how a data-driven application will react in different scenarios and to provide better initialisation and testing for data-driven applications[22]. A *valid knowledge-driven model can be used as source of data* that can be used to train, test and enhance the predictive power of data-driven models[23,24]. Alternatively, explicitly including the knowledge-driven model as prior knowledge will help in *reducing the training set* required to train the data-driven models. Several data-driven methods have a way of explicitly accounting for causal prior knowledge in the derivation of the predictive model such as Physics Informed Neural Nets[25]. This significantly reduces the amount of data needed to train the data-driven model and improves its predictive capability. Additional work is required as these methods are currently computationally quite slow, however this is a field in rapid evolution with fast progress. Finally, *transfer learning* can be used, where a data-driven model is initialised by first training against data generated from a knowledge-driven model and is then re-trained against a reduced but tractable real data set to tune the model[26]. These approaches of combining knowledge-driven virtual physiological human models with data-driven models have now been adopted in clinical applications, with data-driven models being trained on virtual patient cohorts[27,28,29,30] (**example 3**, p. 19).

Another way knowledge-driven models (or the in data they generate) can be used to enhance data-driven models is by providing a *benchmark against which the numerical accuracy of the developed algorithm itself can be tested* for model verification (not validation) purposes. Such benchmarks are already relatively common in knowledge-driven modelling, when the level of model complexity requires the verification of off-the-shelf dedicated solvers, *e.g.,* in computational coupling of electro-mechanics phenomena in physiology[31]. Similar verification processes can be applied for the verification of the capacity of data-driven models to duly learn expected basic relationships between different inputs and outputs spaces.

22: Kalra N and Paddock SM, Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? RAND Corporation, 2016. https://www.rand.org/pubs/research_reports/RR1478.html.

23: Savva M, et al. Habitat: A Platform for Embodied AI Research; Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2019, 9339-9347.

24: Kolve E, et al. A. Ai2-thor: An interactive 3d environment for visual ai. arXiv 2019, 1712.05474v3 https://doi.org/10.48550/arXiv.1712.05474.

25: Karpatne A, et al.. Physics-guided neural networks (pgnn): An application in lake temperature modeling arXiv 2021, 1710.11431v3. https://doi.org/10.48550/arXiv.1710.11431.

26: Zappone A., et al. Wireless Networks Design in the Era of Deep Learning: Model-Based, AI-Based, or Both? IEEE Transactions on Communications, 2019;67(10):7331-7376. doi: 10.1109/TCOMM.2019.2924010.

27: Giffard-Roisin S, et al., Transfer Learning From Simulations on a Reference Anatomy for ECGI in Personalized Cardiac Resynchronization Therapy. IEEE Trans Biomed Eng. 2019;66(2):343-353. doi: 10.1109/TBME.2018.2839713.

28: Alawad M, Wang L. Learning Domain Shift in Simulated and Clinical Data: Localizing the Origin of Ventricular Activation From 12-Lead Electrocardiograms. IEEE Trans Med Imaging. 2019;38(5):1172-1184. doi: 10.1109/TMI.2018.2880092.

29: Niederer SA, et al., Creation and application of virtual patient cohorts of heart models. Philos Trans A Math Phys Eng Sci. 2020;378(2173):20190558. doi: 10.1098/rsta.2019.0558.

30: Roney CH, et al.,. Predicting Atrial Fibrillation Recurrence by Combining Population Data and Virtual Cohorts of Patient-Specific Left Atrial Models. Circ Arrhythm Electrophysiol. 2022;15(2):e010253. doi: 10.1161/CIRCEP.121.010253..

31: Lluch È, et al., Breaking the state of the heart: meshless model for cardiac mechanics. Biomech Model Mechanobiol. 2019;18(6):1549-1561. doi: 10.1007/s10237-019-01175-9.

8. Credibility assessment of knowledge-driven and data-driven models in health care

Several guidance documents and/or standards outline a procedure for **assessing credibility** of knowledge-driven models, defined in the context of medical devices and medicinal products[12,32,33]. After identification of the question of interest, describing the problem the model shall address, the CoU is defined, describing the role and scope of the model and how it is going to be used in relation to other forms of evidence. Then, the model risk is assessed for the specific CoU, taking into account the influence of the model on the medical decision as well as the consequence to the patient of an incorrect decision. These elements combined allow to set credibility goals for the model which will be obtained through verification, validation and applicability analysis.

This strategy for credibility assessment follows a "fit-for-purpose" approach (all models are wrong, but some are useful). It acknowledges that no model will be a perfect picture of reality but - still - it may serve for a certain CoU and associated risk given that it fulfils a certain set of quality criteria. This means that, even though the same model might be applicable, it must not be used in a different context (and associated risk) where the credibility would require more stringent quality criteria, without re-iterating the credibility.

For closed data-driven models (models that do not change after they have been released for use), the credibility assessment can in many ways (though not entirely) be similar to that of the knowledge-driven models mentioned above. Several AI models have already obtained permission from the FDA for use in clinics[34].

However, **self-learning models** do pose a substantial credibility assessment challenge as every new data set that is entered has the potential to change the entire AI model, which means it would require a complete re-validation. How to address this challenge is the subject of current regulatory discussion[35] with FDA considering a lifecycle-based regulatory framework for these technologies that would allow for modifications to be made from real-world learning and adaptation, while still ensuring maintenance of the safety and effectiveness of the software as a medical device.

One additional element to mention here is the **explainability of AI**, referring to the ability to understand how the algorithm reached a certain solution. Explainability is explicitly requested by the EU-GDPR[36] (article 15 mentions the right of explainability and article 22 the right of human intervention). Many of the more modern AI algorithms, such as deep learning, are naturally opaque (*i.e.* not explainable) but additional methods are being developed to address this issue[37]. Knowledge-driven models are explainable by their very nature.

> ***"No model will be a perfect picture of reality but - still - it may serve for a certain context of use and associated risk given that it fulfils a certain set of quality criteria.***

32: Kuemmel C, et al. Consideration of a Credibility Assessment Framework in Model-Informed Drug Development: Potential Application to Physiologically-Based Pharmacokinetic Modeling and Simulation. CPT Pharmacometrics Syst Pharmacol. 2020 Jan;9(1):21-28. doi: 10.1002/psp4.12479.

33: Musuamba FT, et al., Scientific and regulatory evaluation of mechanistic *in silico* drug and disease models in drug development: Building model credibility. CPT Pharmacometrics Syst Pharmacol. 2021 Aug;10(8):804-825. doi: 10.1002/psp4.12669.

34: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices?utm_source=FDALinkedin#resources

35: https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device

36: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02016R0679-20160504&from=EN

37: Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021;3(11):e745-e750. doi: 10.1016/S2589-7500(21)00208-9. PMID: 34711379.

9. Conclusion and policy asks

Health policies can lead to facilitating the development of powerful *in silico* technologies (in the widest meaning as discussed above). The explicit inclusion of modelling in regulatory legislation (Medical Device Regulation (MDR) and *In Vitro* Diagnostics Regulation (IVDR)) has given an impetus to the relevant stakeholders to consider modelling and simulation in their respective processes. The European Commission's creation of the European Health Data Space[38] and the proposed AI legislation[39] will have an important impact on the further development of not only data-driven models but of *in silico* medicine as a whole. Reciprocally, *in silico* models, be they knowledge-driven, data-driven or hybrid, can drive health policies. Given the European Commission's Pharmaceutical Strategy3[40] (including the revision of the orphan and paediatric regulations), *in silico* models can have a major impact on the risk assessment and overall safety of medicinal treatments and contribute to making more affordable new drugs (and vaccines) available by reducing their development time and cost.

To date, EU health policy does not put forward any formal regulatory frameworks for the various applications of *in silico* medicine, and as such it remains a grey area when it comes to legislative guidance. In order to ensure the systematic use of *in silico* medicine, it is crucial to strive for a policy framework that can (among others) enable *in silico* medicine in a manner that can raise safety standards in clinical trials, improve the lives of rare disease patients and reduce costs to device manufacturers and researchers. Stakeholders require legal certainty from regulators on what constitutes an acceptable model and as such need a framework which they can rely on to employ these techniques for safer and more effective healthcare delivery. This framework should follow existing EU precedents which give room to manufacturers to perform conformity assessment by showing compliance with EU legislation through harmonised standards, or by other means of their own choice (*e.g.*, by means of any existing technical specifications including all other available international standards).

In line with this context, the members of the Avicenna Alliance from medtech, pharmaceutical, software and life science industries, as well as from the in silico medicine academic community and healthcare organisations call on the European Commission to:

- Reflect the added value that *in silico* technologies have in healthcare within the revision of existing pharmaceutical legislation. This is particularly the case for the production of digital evidence for treatment development, especially for paediatric and rare diseases, which cannot be generated by the *in vitro* or *ex vivo* models, or is very challenging to generate using *in vivo* (animal and patient) models.
- Introduce a clear regulatory framework for digital evidence for all medical products, including both drugs and devices, that takes into account the context of use and model risk to reduce political and economic barriers.
- Develop a Good Simulation Practice, similar to Good Clinical Practice or Good Manufacturing Practice, as a starting point for developing *in silico* models used in decision-making processes such as regulation.
- Launch funding schemes promoting the adoption of *in silico* technologies in healthcare through pre-commercial procurements, public procurements of innovations and early dialogues either with local/national health technology assessment agencies or through EU-level joint clinical assessments.
- Promote public trust in *in silico* technologies by increasing awareness, e.g. increasing public literacy about said technologies.

38: https://ec.europa.eu/health/ehealth/dataspace_en

39: https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence

40: https://ec.europa.eu/health/human-use/strategy_en

Example 1

Hypertrophic cardiomyopathy (HCM) is a disease in which the heart muscle becomes abnormally thick. Electrocardiogram (ECG) is one of the simplest and fastest tests to evaluate the heart. Although there are general links between changes in the ECG and changes in the heart morphology and sudden cardiac death, the coupling between HCM and ECG remains poorly understood with suboptimal patient stratification as a consequence. In this example, data-driven modelling (statistical methods, blue part) is used to **classify the patients** into specific subgroups based on the similarities in their ECG signals [41,42,43].

Subsequently, knowledge-driven modelling was used (red part) starting from the patient's medical images of the heart to **link the identified subgroups to abnormalities in underlying biological processes** such as problems in the ion channels that are important for the production of electrical signals in the heart (leading to specific ECG signals). The combination of both data-driven and knowledge-driven approaches allowed to identify personalised treatment strategies based on ECG read-outs.

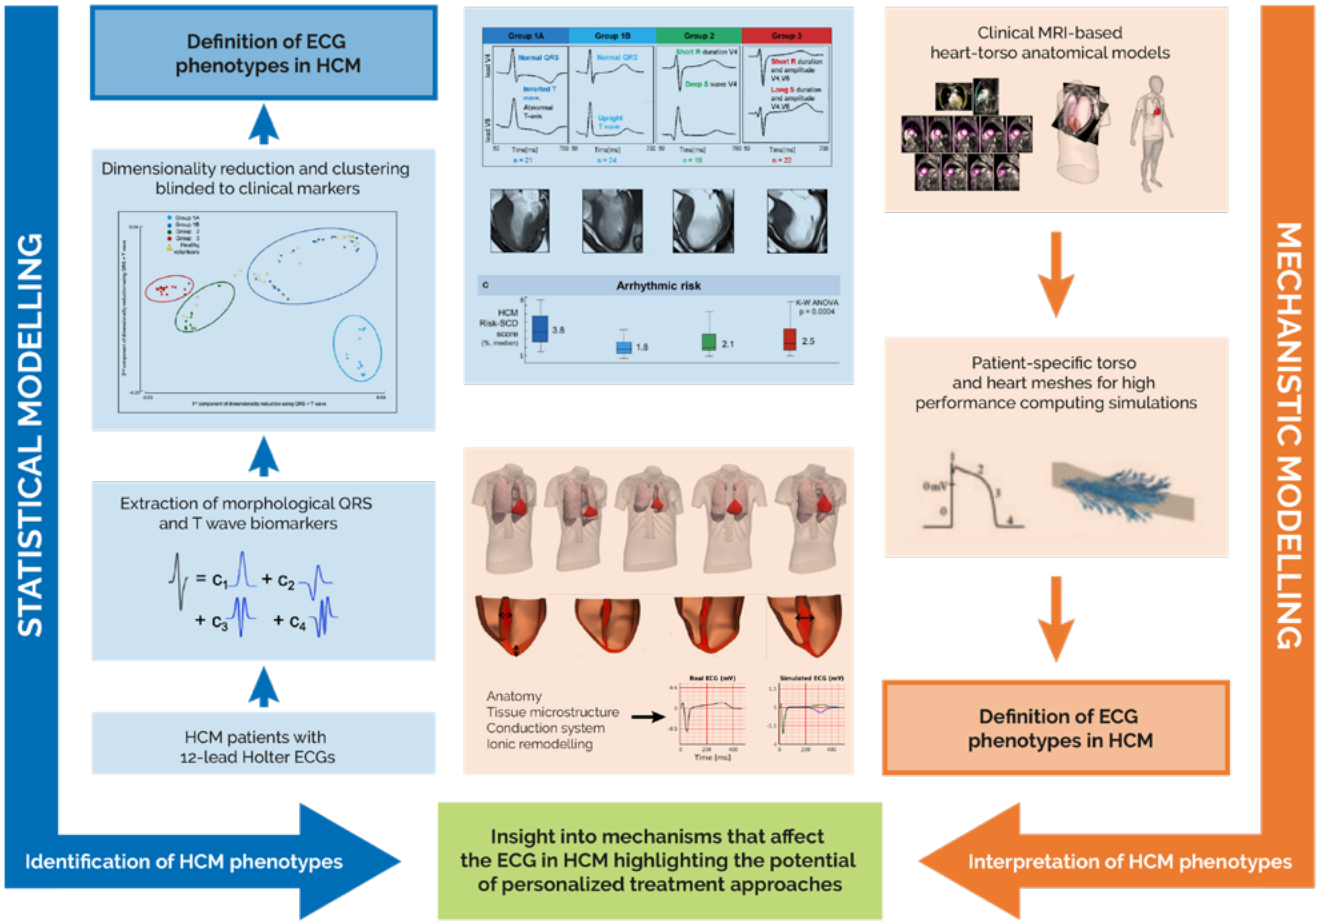

Figure 2:
Synergy between knowledge driven (mechanistic) models and data driven (statistical) models in the definition of electrocardiogram (ECG) biomarkers for the management of hypertrophic cardiomyopathy (HCM). Reproduced from Corral-Acero et al. (2020)[21], with permission.

41: Lyon A, et al., Electrocardiogram phenotypes in hypertrophic cardiomyopathy caused by distinct mechanisms: apico-basal repolarization gradients vs. Purkinje-myocardial Coupling abnormalities. Europace 2018;20: III102–III112.

42: Lyon A, et al., Distinct ECG phenotypes identified in hypertrophic cardiomyopathy using machine learning associate with arrhythmic risk markers. Front Physiol 2018;9:213.

43: Corral-Acero J, et al., The 'Digital Twin' to enable the vision of precision cardiology. Eur Heart J. 2020 Mar 4:ehaa159. doi: 10.1093/eurheartj/ehaa159.

Example 2

Musculoskeletal models can be used to explore the mechanical balance of the human body and support efficient physical rehabilitation therapies. They can also be used for holistic assessment of patients with multifactorial joint disorders, such as osteoarthritis or juvenile rheumatoid arthritis, where mechanical loads play an important role. Using gait analysis (motion capture) and specific computational techniques, the forces acting on the joints can be calculated. Translating these forces into loads experienced by the tissues of the joint such as the cartilage requires computationally intensive (expensive) simulations. That is not an option if these models are to be used in clinical evaluations. Thereto, a **surrogate model** is built.

First, the knowledge-driven model is simulated for a wide range of motions and material properties. Subsequently, the obtained (modelling) data is used by a data-driven technology to link the joint loads to the modelled output. The resulting artificial neural network can be solved in real time, allowing the clinician to evaluate the loading in the tissues by performing gait analysis[44].

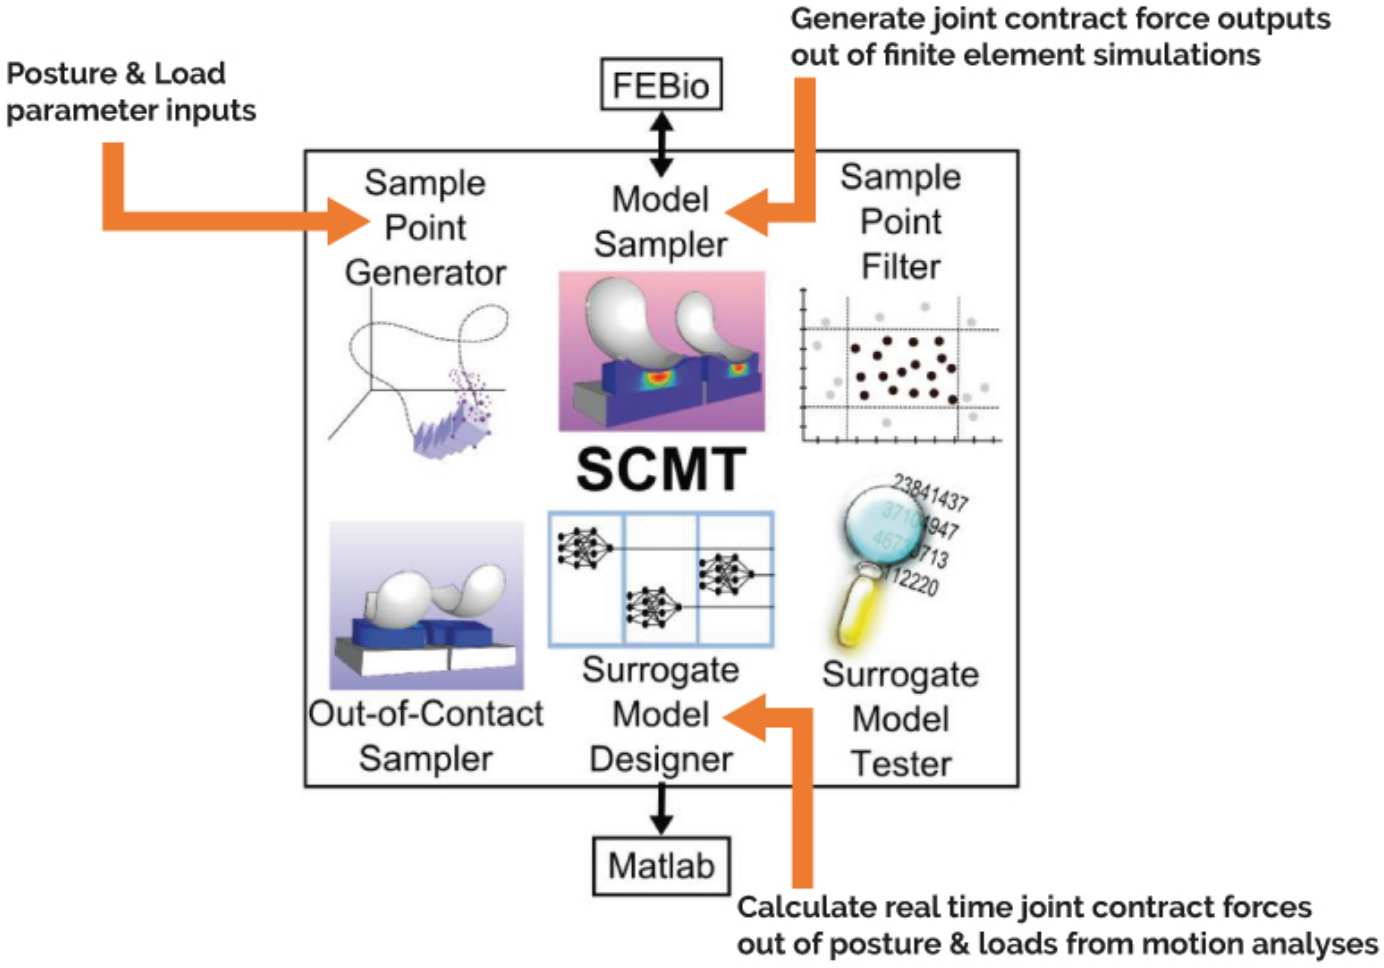

Figure 3:
The Surrogate Contact Modelling Toolbox (SCMT) composed of several tools or modules. Adapted from Eskinazi & Fregly (2016)[23], with permission.

44: Eskinazi I, Fregly BJ. An Open-Source Toolbox for Surrogate Modeling of Joint Contact Mechanics. IEEE Trans Biomed Eng. 2016;63(2):269-77. doi: 10.1109/TBME.2015.2455510.

Example 3

This example demonstrates the use of a **machine learning model trained using virtual patient cohorts** to predict the optimal ablation procedure strategy in patients with atrial fibrillation[28]. In patients with atrial fibrillation, an irregular activation of the top two chambers of the heart, ablation can be used to remove or isolate tissue that is sustaining the arrhythmia.

However, only a single ablation strategy can be delivered to each patient, this means it is not possible to determine if the optimal strategy, that treats the arrhythmia but minimises the amount of tissue ablated, was delivered in any single patient. The acute outcome of multiple ablation strategies can be delivered in virtual patients. By creating a virtual patient cohort and delivering multiple ablation strategies to each patient, it is possible to train a machine learning model to predict which ablation strategy will be best for each patient and identify the factors that best inform this decision. This creates a platform for training data-driven models when the clinical data cannot be collected.

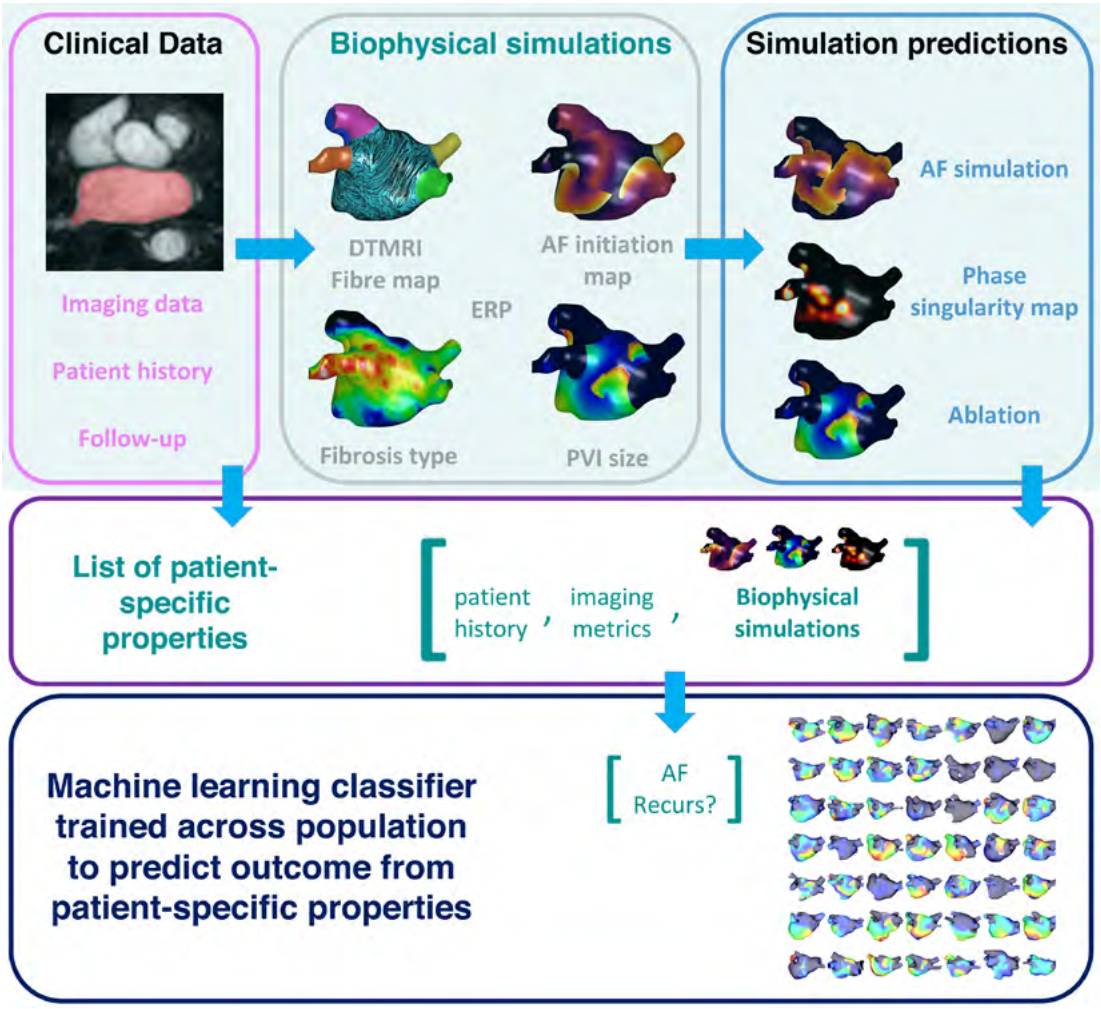

Figure 4:
Schematic methodology for using machine learning to combine biophysical simulation stress tests for acute simulation responses with population data to predict long-term atrial fibrillation (AF) recurrence. DTMRI: diffusion-tensor magnetic resonance imaging; ERP: effective refractory period; PVI: pulmonary vein isolation. Reproduced from Roney et al. (2022)[28] with permission.

About the Avicenna Alliance

The Avicenna Alliance is an association of industry and research organisations who have a commercial or research interest in the development of In silico medicine. The Association, established in 2015, has its origins in the Virtual Physiological Human Initiative, a European Commission funded scientific domain focused on research into computer modelling and simulation. Tasked by the European Commission with developing a **"Roadmap for In silico medicine"**, the Association now seeks to put this roadmap into policy and ensure the development of a regulated In silico market.

This Association bridges the gap between the scientific community, industry and policy makers by advocating for policy changes that take into account scientific and market developments.

About the Virtual Physiological Human institute

The **Virtual Physiological Human Institute**, in short VPH Institute, is an international non-profit organisation, whose mission is to ensure that the Virtual Physiological Human is fully realised, universally adopted, and effectively used both in research and clinic. To this end, it organises its activities both on the scientific level and the policy-regulatory level, ensuring the entire path from computer screen to the patient is rolled out and all stakeholders are involved.

Credits

Published August 2022, © Avicenna Alliance and VPH Institute.

Graphic Design: Giulia Bassetti

Text revision: Liesbet Geris

This paper was written by members of the AI Task Force of the Research & Technology Working Group of the Avicenna Alliance. We gratefully acknowledge input from Dr. Prasanna Pathmanathan (FDA, CDRH).