European Review for Medical and Pharmacological Sciences 2020; 24: 7462-7474

Artificial intelligence: a disruptive tool for a smarter medicine

C.M. GALMARINI, M. LUCIUS

Topazium Artificial Intelligence, Madrid, Spain

Abstract. – OBJECTIVE: Although highly successful, the medical R&D model is failing at improving people's health due to a series of flaws and defects inherent to the model itself. A new collective intelligence, incorporating human and artificial intelligence (AI) could overcome these obstacles. Because AI will play a key role in this new collective intelligence, it is necessary that those involved in healthcare have a general knowledge of how these technologies work. With this comprehensive review, we intend to provide it.

MATERIALS AND METHODS: A broad-ranging search has been undertaken on institutional and non-institutional websites in order to identify relevant papers, comments and reports.

RESULTS: We firstly describe the flaws and defects of the current R&D biomedical model and how the generation of a new collective intelligence will result in a better and wiser medicine through a truly personalized and holistic approach. We, then, discuss the new forms of data collection and data processing and the different types of artificial learning and their specific algorithms. Finally, we review the current uses and applications of AI in the biomedical field and how these can be expanded, as well as the limitations and challenges of applying these new technologies in the medical field.

CONCLUSIONS: This colossal common effort based on a new collective intelligence will exponentially improve the quality of medical research, resulting in a radical change for the better in the healthcare model. AI, without replacing us, is here to help us achieve the ambitious goal set by the WHO in the Alma Ata declaration of 1978: "Health for All".

Key Words:
Artificial intelligence, Computational intelligence, Computer vision systems, Biomedical research, Oncology.

Introduction

Around 300 billion US dollars are invested every year in medical research worldwide. The result of this collective effort has been extraordinary: in just 100 years, life expectancy has rocketed from approximately 45 to 72 years. However, there is still a lot of work to be done since the figure for all-cause mortality stands at almost 57 million. This figure includes 47 millions of adults with chronic, non-communicable diseases, and 6.5 millions of children, of whom 5.6 million are under the age of five[1]. In spite of the massive economic efforts in R&D and innovation, global health improvement seems to have reached a plateau. The current model of medical research is showing difficulties due to a series of limitations[2]. Firstly, its current framework mirrors a linear and sequential process, which has been the only paradigm accepted so far. In this linear approach, investment in R&D produces results that are, at best, only proportional to the level of investment. Furthermore, in the current R&D model the outcome information is usually stored in closed and sealed compartments with very limited access to outsiders, sometimes even within companies or institutions. These "Chinese walls" not only prevent innovation, they also increase the time and cost of delivering new drugs or medical devices to patients. Additionally, the current model assesses "value" as the amount of "assets" an organization owns (patents, chemical entities, etc.) and the relevance of an organization in the market is measured based on these assets. However, the paucity of data encourages monopolization of resources and discourages the sharing of information, which in turn hinders innovation. Thus, the current model drives pharmaceutical and biotechnology industries to generate innovation and value through the acquisition of assets and discoveries from other stakeholders, which results in overcrowding of companies within this sector. All these hurdles have led to a R&D model that no longer works as fast as wanted.

A New Medical R&D Model Is Needed

Investing more financial and human resources is not enough to modify the current R&D model

Corresponding Author: Carlos M. Galmarini MD, Ph.D; e-mail: cmgalmarini@topazium.com

in a way that would allow to achieve the ultimate goal of improving healthcare. A radical change is needed. This should be based on a new form of collective intelligence able to overcome the obstacles of the current model. Since it appears that human intelligence is not enough, this new type of collective intelligence should include artificial intelligence (AI)[3]. We can use this new kind of collective intelligence to create smarter organizations, conduct smarter research and solve medical problems that, up to now, where thought to have no solution[4]. How and why will AI revolutionize medical research? It is known that when a process becomes digitized and powered by information flow, its rate of development jumps onto an exponential growth path, doubling every two years (Moore's Law). Indeed, this system encourages *in silico* research, which exponentially reduces the research costs and promotes innovation, resulting in a wealth of information and a further cost reduction, thus generating a positive feedback loop where the output dramatically increases, since all unnecessary processes and expenditure on purposes other than the generation of value have been eliminated. This information-led environment represents a paradigm shift, as it promotes knowledge sharing and the discovery of innovative solutions. In this new model, the most benefited are those who share more knowledge and acquire new knowledge faster.

However, AI elicits many fears, as it is thought that it has the potential to replace human intelligence. Yet these fears are unfounded, for AI should be viewed as a tool rather than a replacement. Therefore, what can AI add to human intelligence? Nowadays, we are overwhelmed with information from different new sources (e.g., mobile phone technology, wearable devices, etc.). AI can integrate these vast amounts of data and identify relevant patterns in a way that a human mind would be unable to do. On the other hand, AI cannot replace human judgment and wisdom. Therefore, human and machine intelligence can complement each other[5]. Because AI will play a key role in medical research, it is necessary that those involved in this field have a general knowledge of how these technologies work. Thus, the aim of this review is to offer a comprehensive background on AI, as well as its current uses and applications in the medical field, and how these can be expanded in the future.

Almost Human: Artificial Intelligence

AI is the field of computer science dealing with the development of computers that can act "intelligently"[4]. Top applications in healthcare include drug discovery and development, automated devices, wearables, diagnostic and medical imaging devices, remote monitoring of patients, predictive medicine, robotics and management of healthcare systems[6]. There are two key factors that allowed progress in the development of AI (Figure 1). Firstly, in recent years, new sources of medical data, such as health monitoring devices, mobile apps or electronic health records (EHRs) have been developed, leading to the generation of vast data banks containing patient data. A second factor is the access to new analytical tools that can mine and process vast amounts of data. This allows to perform pattern recognition, data classification and predictions in a much shorter period of time.

Data Access

The success of AI relies on its ability to access the right kind of data in an appropriate volume. This represents an enormous challenge. It is calculated that approximately a zettabyte (a trillion gigabytes) of medical data is generated each year. Due to the existence of different technologies that allow to collect data more easily, this amount is expected to double every two years[7,8].

New Forms of Data Collection

Wearables can deliver an unprecedented amount of data from millions of people. For example, smartphones are one of the most powerful devices for health monitoring. Indeed, these regular devices contain sensors, as well as processing and communication capabilities that allow data to be collected, analysed, and shared. Collecting daily measurements with a smartphone may reveal nuances that are not evident in monthly clinic visits. Similarly, in recent years, increased attention has been given to patient-reported outcomes (PROs), which provide a direct indicator of a patient's health condition without correction or interpretation by a clinician[9-12]. These data are complementary to those collected by the doctor as they provide additional information on other aspects that are important to the patient and that, perhaps, clinicians do not expect or take into account. PROs thus are essential components in high-quality, patient-centred care[13,14]. Collection and integration of ePRO into routine care is feasible, as demonstrated by different studies showing how it increases clinician awareness, resulting in an early response to patient symptoms and ultimately, a better quality of life[14,15].

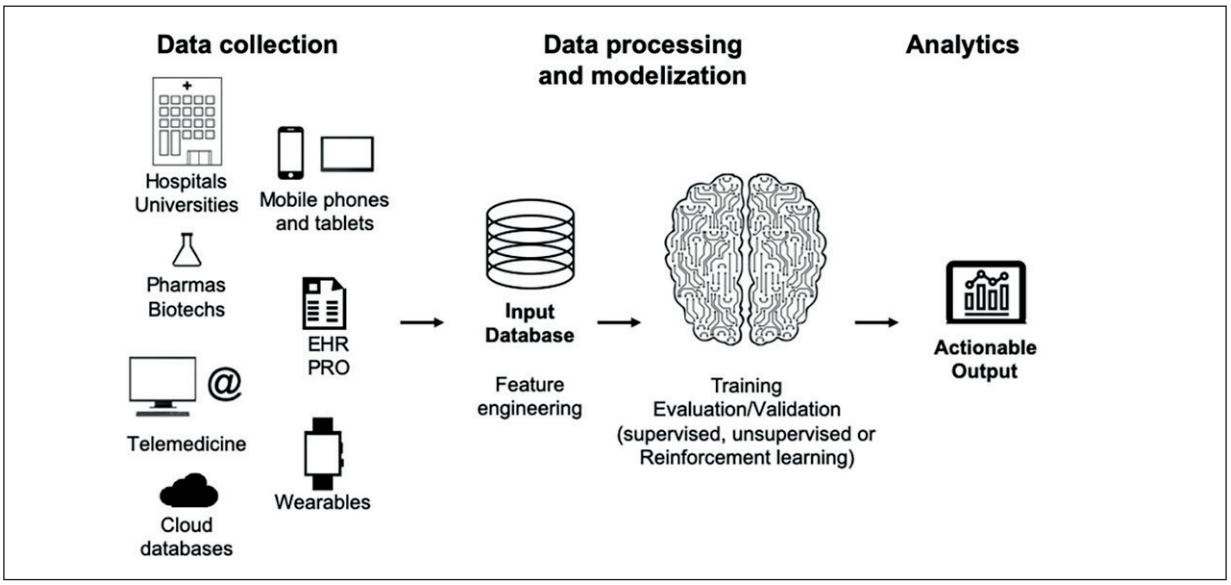

Figure 1. Overview of the AI workflow. The first step is data collection. Data can be collected from different devices and databases generated by universities, hospitals, and biotech and pharmaceutical companies. The second step comprises data processing and data modelling, including feature engineering, model training, evaluation and validation. During this stage, the system can process vast amounts of data and deliver an actionable output that can be a pattern, a classification and/or a prediction.

Data Processing

The completeness and quality of the collected data will determine the actual usefulness of the database, since the performance of any model relies on the data used to train it[16]. Thus, formatting, filtering, and normalizing the collected data is a first critical step, and represent most of the effort in building any analytical model[17,18]. Many different challenges arise when building a useful database. A first problem is that the inputted dataset may be incomplete. The options for dealing with missing data include directly inferring the missing values or simply removing sparse features. Another challenge is that there can be thousands of potential predictor features in EHRs[19]. Traditionally, this problem was tackled by simply considering only a very limited number of commonly collected variables[20]. However, this poses a problem, since the resulting models may lead to inaccurate predictions[21]. Moreover, not every input feature in a given dataset will be relevant for predicting the output label. In fact, including irrelevant input features can compromise the performance of the model derived by AI. A process called feature selection is often used to identify relevant input features. An example of a feature selection technique is to correlate all input features with the output labels: only those features that meet a pre-defined threshold of relevance are kept. However, in healthcare there is an additional hurdle: patient data must be anonymized and pooled. New developments, such as blockchain could be used to build data systems that allow people to easily and securely share their health data with researchers, while retaining control over the information supplied[22]. A second important step when preparing data for an AI model is transforming the input data into a language that a computer can understand (a process known as encoding). Encoding must be done in such a way that the semantic characteristics of the data are captured. Usually, data and categorical and scalar values are represented as binary vectors. Numerical features are typically zero-centered by a subtraction of the mean value from every data point or normalized to variance. Finally, log transformations are used when extreme values are altering a feature distribution.

Over the last decade, we have witnessed a dramatic increase in the number of large, highly complex databases being generated from wide-ranging datasets, including human genome sequencing, gene expression profiles, proteomics, metabolomics, microbiome, high-resolution imaging, EHR, etc.[23]. Within a few years, we will have vast comprehensive information from millions of patients[24]. One key challenge would be encoding and combining this diversity of data, stored in multiple sites across many different organizations within the health sector. Platforms are interfaces that allow

people to connect and share data and insights rapidly, simply and securely. In the healthcare arena, these platforms can help to connect very unrelated stakeholders (e.g., patients, practitioners, researchers and regulators), who can combine capabilities and share data to overcome the limitations of the current sequential approach. AI and data analytics combined have the potential to become "platform aggregators".

Algorithms

AI involves different applications of advanced computer intelligence. Machine learning is a subfield of AI that includes a number of data analytics techniques aimed at building predictive models from multidimensional datasets; in other words, computers identify patterns from inputted data without being programmed[16]. The canonical machine learning workflow involves processing the input data, training the underlying model (which consists of a set of mathematical formulas and statistical assumptions defined by the learned rules), and then, using it to predict new data. The data inputted to a machine-learning algorithm typically consist of "features" and "labels" across a set of samples. Features are the measurements or characteristics, either raw or mathematically transformed, are used to classify the samples, while labels are the outputs that the algorithmic framework intends to predict. Features and labels can be continuous or categorical. Data may require a primary pre-processing. During this stage, missing or spurious elements are identified and addressed while the remaining data are transformed into a format more adapted for an algorithm. This process is called featurization or feature engineering. The more appropriate the representation of the input data, the more precisely an algorithm will map it to the output data[25].

Depending on the type of data and the question to be answered, a variety of machine learning algorithms can be applied. Naive Bayes classifiers are classification algorithms based on the Bayes' theorem that, given a set of existing data, calculates the probability of a hypothesis (or model) being true[26]. On the other hand, k-nearest-neighbour methods calculate the distances between the test sample and the samples in the training data[27,28]. Nearest neighbour methods can be utilized for both classification and regression. Another approach is decision trees, which are flowchart-like diagrams used to determine an outcome[29]. Decision trees are often used in meta-algorithms or ensemble methods, which combine multiple trees into one predictive model to improve performance. Furthermore, kernel methods are a class of algorithms, amongst which support vector machine and kernel ridge regression are the most familiar[30]. The name 'kernel' is due to the use of a kernel or similarity function, that maps non-linear input data into a higher-dimensional space allowing its processing by linear classifiers. Finally, and differently from the other machine learning methods that focus mostly on pattern recognition, reinforced learning is oriented to experience-driven sequential decision-making. These algorithmic frameworks are capable of taking actions in a particular environment to maximize cumulative reward. Since output labels can be continuous or categorical, many machine-learning tasks include regression or classification, respectively, where a regression task consists in the prediction of continuous output variables and a classification task consists in the prediction of categorical output variables.

Once the algorithm or set of algorithms have been selected, the trial model must be judged to allow its optimization. The learning process itself consists in finding the most relevant parameters for the model, i.e., those that make it possible to translate input data into accurate output data. These parameters are identified through a series of back-and-forth iterations, where parameters are tested, the performance of the model is evaluated, errors are identified and corrected, and then, the process is repeated again. This process continues until the performance of the model can no longer be improved, as determined by minimization of the model error. Once the most relevant parameters have been identified, the system can be tested by making predictions for new input data. If the model is accurate on both the training and the test data, then, we can say that the model has been properly trained. Machine-learning algorithms can be trained using supervised, unsupervised, semi-supervised or reinforced learning approaches[31,32]. Supervised learning is applied when there are labels available for the input data[16]. In this case, the training data consist of sets of input features and associated output labels that are related. The labels are used to train the machine-learning framework through identification of patterns that can predict the output labels. The goal of the algorithm is to originate a function that, given a specific set of input values, predicts the output values to an acceptable degree of accuracy. Subsequently, when new input data become available, the predictions using the

derived model can be directly made. Supervised learning includes classification algorithms (for categorical output variables; e.g., support vector machine) or regression algorithms (for continuous output variables; e.g., linear and logistic regression). By contrast, unsupervised approaches are utilized when the output labels for the input data are not known; these methods learn only from the patterns in the input data, i.e., they identify factors about input features without known output labels. Thus, the goal of unsupervised approaches is to cluster subsets of the data based on similar features and to identify the number of clusters present in the data. Unsupervised learning cannot produce an independent predictive model. Commonly used unsupervised methods include clustering algorithms (e.g., hierarchical clustering) and dimensionality reduction algorithms (e.g., principal component analysis). Unsupervised methods can be advantageous when output labels are missing or incorrect, as it can still identify patterns, since the clustering is performed only on the input data. Of note, the output obtained by models built through unsupervised approaches can be used as input data for models built by supervised approaches. This is known as semi-supervised learning, which may be useful if there is a large amount of input data without the corresponding labels. The algorithm learns from labelled data to make a prediction for unlabelled data and identify patterns. Finally, reinforcement learning involves a set of dynamic algorithmic frameworks that continuously learns from the environment in an iterative fashion using a system of reward and punishment. This type of learning is limited to a particular context because what may produce a maximum reward in one situation may be directly associated with punishment in another. Algorithms learn to predict the flawless behaviour within a specific context, to maximize its performance and recompense. Thus, learning occurs *via* a reward feedback, known as reinforcement signal.

A machine-learning method is selected depending on the task, the characteristics of the data, and the labelled or unlabelled nature of the data. As stated above, if the data are labelled, then, a supervised method can be applied to generate a predictive model by either regression or classification of the data. If the data are unlabelled, then, learning must be conducted by an unsupervised approach. After building a model by the appropriate machine learning method, the outputs must be validated. New data can be generated, collected and used to refine the derived model, improve its accuracy, and develop new hypotheses. However, the reliability of machine-learning methods can be compromised by various elements[27]. Excessive variance ("overfitting") or excessive bias ("underfitting") is always associated to a deficient performance of a model. Overfitting appears when the parameters of the training data are fit so specifically to the model that they do not have predictive power outside this particular dataset; typically, this occurs when the number of parameters is important, and the model becomes too complex. The issue of overfitting can be addressed by expanding the size of the training set or by decreasing the model complexity. On the other hand, underfitting occurs when the model cannot detect an accurate relationship between input and output data, or when the available data are inadequate to permit an accurate pattern identification. Underfitting can be addressed by increasing the model's complexity[33]. The main way to test the accuracy of a machine-learning model is to check whether it can successfully predict unknown output data. This is usually done by setting aside a randomly selected portion of the dataset during the model's training. When the training is finished, this set-aside data (also called validation dataset) are inputted to the model. The degree to which the output data in the validation set is accurately predicted by the derived model provides a measure of its quality.

Deep Neural Networks

Deep Neural Networks (DNNs) is one of the principal algorithmic frameworks in machine learning. These are algorithms based on interconnected artificial neurons that execute transformations on non-linear data to identify relevant predictive features based on examples[34,35]. DNNs are excellent gears for identifying patterns in extremely large and complex datasets. A classical DNN architecture include a number of layers. The first ones mine relevant features from the input data, and the last ones correlate the extracted features to outputs. The width of a DNN relates to the maximum number of artificial neurons in a layer while its depth, to the number of hidden layers it comprises. A neuron (also known as a node or perceptron) is a computational unit with one or more connections receiving the inputs, a transfer function associating the inputs, and an output connection. The neurons in the input layer take the raw data and transfer it to the neurons in the hidden layers, where computation happens. The

hidden layers are sequentially connected such that each of them learns properties about the data by taking as input the transformed input or representation produced by the previous hidden layer. Once this representation is fed to the next layer, this is transformed into a new representation. Indeed, in each neuron, input data are combined with a set of coefficients that amplify or quench it so, assigning weights to inputs within the context of the model being built. These weighted inputs are summed, and the result is run through a non-linear activation function, to conclude whether that signal should progress any further through the DNN. If the signal is transferred, is said that the neuron has been "activated". Max-Out, Sigmoid, rectified linear unit (ReLU) or Softmax are some of the most common activation functions. Depending on the activation function being used, the properties of the activation for the DNN can be quite different. Errors originated from the training labels are back propagated across the DNN and the model tuned to achieve a higher performance. Learning is the process of tuning the weights so that the training data are represented as exactly as possible. Finally, the output layer generates a prediction based on the weighted inputs from hidden layers[25]. Once trained, DNNs can be used over new datasets to accurately generate output values. DNNs provide considerable advances compared to conventional machine learning methods. Indeed, DNNs can be utilized to attack extremely complex problems and identify patterns in datasets that are just too large and complex for a human brain to process. Similar to conventional machine learning, DNNs can be used with supervised, unsupervised, semi-supervised or reinforced learning methods[36].

DNNs can have a variety of architectures. The simplest of all has three layers: an input, a middle (hidden) and an output (prediction) layer. Common DNN architectures often incorporate additional functions and connections to the typical architecture to permit training using deeper networks. Nowadays, DNNs can be simply build from pre-existing modules, such as Caffe, Theano, PyTorch and TensorFlow. These frameworks differ principally in easiness of use and manipulability. One of the principal challenges in deep learning is how to select the proper DNN architecture for a given task. The ideal architecture should be able to achieve truthful results with the fewest parameters. Therefore, the number and size of the hidden layers needs to be defined based on the purpose of the model to be trained. Convolutional Neural Networks (CNNs) are one of the most common architectures used for deep learning[37]. They are normally comprised of numerous convolutional and pooling layers. This structure favours learning abstract features at increasing scales going from object edges up to entire objects. That is why CNNs are well suited for image recognition and computer vision tasks[38,39]. CNNs use weight sharing by sliding a tiny, trainable filter of weights across the input vector and convoluting each overlapped region of inputs with the filter. This differs from fully connected neural networks where activation units are bound to all inputs of a feature vector. For example, CNN's "hidden layers" frequently comprise a series of convolution operations that extricate feature maps from the input image and pooling actions that perform feature aggregation. The "hidden" layers output is then input into "fully connected" layers that delivers high-level analysis. Finally, the "output" layer produces predictions. Labelled data and thereby supervised learning methods are frequently used to train CNNs end-to-end.

There are other commonly used DNN architectures. Autoencoders are DNNs that learn how to efficiently compress and encode data, and then, reconstruct it into a representation that is as near as possible to the original input. These types of algorithms have been applied, for example, to the classification of cancer cases using gene expression profiles or for predicting the sequence of proteins. Restricted Boltzmann machine learning uses a stochastic recurrent neural network composed of two layers, one made of visible units and another made of hidden units, with no lateral connections. This is one of the most relatively easy to train DNNs, and allows forming a compact and high-level representation of objects[36]. Restricted Boltzmann machine learning is used for unsupervised pre-training of deep networks previous to the subsequent training of the supervised models, for example, of disordered protein regions or amino acid contacts[40,41]. Deep Belief Networks, used to build probabilistic generative models, are formed by several stacked hidden layers. Each layer is directly connected to the next layer. Recurrent Neural Networks (RNNs) are unsupervised or supervised deep learning architectures for identifying temporo-sequential patterns[42,43]. RNNs can predict the next data point in a sequence using the previous data. RNNs are well suited for analysing protein or DNA se-

quences, as well as clinical databases. Generative adversarial networks (GANs) seek to synthesize new data indistinguishable from the training data. It opposes two neural networks against each other; one called generator that creates synthesized data (e.g., an image), and a discriminator that estimates the probability that a particular sample came from the training data rather than from the generator[44]. GANs have, for example, the potential to speed up preclinical development. Finally, streaming algorithms perform on-the-go calculations[45]. Because data are constantly being inputted at such volume, streaming algorithms try to record the essence of what it has been inputted while strategically ignoring the rest. Streaming algorithms are useful to monitor a database that is constantly being updated. As it happens with other machine learning techniques, deep learning models need to be trained and tested on different datasets to prevent overfitting and ensure that the model can be applied to accurately predict new data.

Applications in Medical Research

AI can be applied to improve a varied range of medical fields (Figure 2). A key advantage is that these new methods can screen volumes of data to identify patterns in complex biological systems that otherwise would be missed.

Drug Discovery and Development

The ever-increasing computational power along with the abundance of extensive datasets have led the scientific community to develop learning algorithms that are capable of shortening time frames for drug discovery and development, and enabling to explore the chemical space more effectively[46-48]. DNNs can be used to improve the prediction of chemical reactions during the synthesis process. Instead, the pharmacological attributes of a molecule can be estimated using atomistic simulations. The generation of different molecules with precise chemical, physical and pharmacological characteristics can be done by training GAN networks while reinforced learning can help the *de novo* design of chemical entities. Also, the establishment of structure-property relationships are estimated by machine learning algorithmic frameworks[49,50]. A large number of computational methods[51,52] for *in silico* drug discovery and target extension have been established. One of these computational approaches is virtual screening, that is, the *in silico* screening of enormous chemical libraries to identify small molecules that bind to a specific protein (drug-target interaction)[53]. Virtual screening can also detect interactions with other non-specific proteins (off-target effects)[54]. Deep learning also allows to effectively assess compound toxicity based on its chemical structure[55]. In biological research, machine learning applications are becoming ubiquitous, encompassing not only genome annotation, but also interpretation of complex biological data comprising drug development, biomarker detection as well as recommendations on clinical targets[56-58]. Another interesting application is the identification of new indications for existing drugs, a process known as drug repositioning or drug repurposing[55,59]. Other areas include the prediction of protein binding sites, identification of key transcriptional cancer drivers or the prediction of metabolic functions in complex microbial communities[60-63].

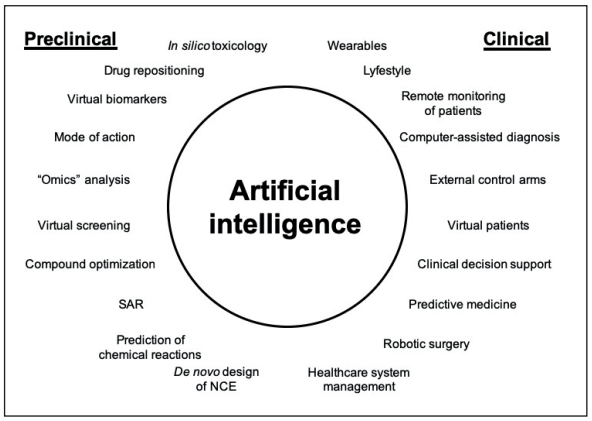

Figure 2. Principal applications of AI in medical research. AI can be applied to improve and accelerate both the preclinical and clinical stages of any project, as well as to guide and support clinical decisions. NCE: new chemical entity; SAR: structure-activity relationship.

Computer-Assisted Diagnosis

Digitalization of medical images (e.g., ultrasounds, computed tomography scans, medical photos, etc.) has promoted the development of automated image interpretation, as almost all image-related tasks involve the evaluation of image features, a field where AI excels. Numerous studies have demonstrated the potential to help radiologists when evaluating a variety of radiological images, including CT scans for pulmonary nodules and pneumonias, mammograms for breast lesions, and magnetic resonance images for brain tumours[64-67]. AI methods can also use any type of image data and therefore,

it can be customized to other fields where images are used, such as pathology, dermatology and ophthalmology. In pathology, AI is emerging as a potent tool to assist pathologists when looking for microscopic lesions in tissue sections[68-70]. Digital histopathology images can be used as inputs for deep learning algorithms specifically proficient, for example, in the discrimination between normal and abnormal tissues[71]. Deep-learning algorithms can also classify skin lesions with a level of competence equivalent to that of dermatologists[72-74]. Similarly, using retinal images (fundus or optical coherence tomography), deep learning algorithms were able to discriminate between macular degeneration and diabetic retinopathy with a performance comparable to that of experts[75,76]. Moreover, deep learning revealed additional traits in retinal images that allowed to better predict cardiovascular risk (e.g., high blood pressure, heart attacks, etc.)[77,78]. These fast, scalable methods can be implemented on mobile devices allowing low-cost universal access to important diagnostic care anywhere, positively affecting primary care practices[72]. All these studies show how to use AI to perform simple, cost-effective, and widely available studies that could help to identify at-risk patients requiring referral to a consultant.

Clinical Applications

AI has also several applications at the clinical level. Input data from a healthy cohort, can be used to train deep learning algorithms on the fundamental characteristics of healthy states. Subsequently, this algorithm could be fed by data from a patient cohort and used to predict differences between healthy and disease states, identifying differences in regulatory interactions and biomolecules that could be validated and explored further. Thanks to AI, physicians would also monitor their patients and check if they are likely to present a subsequent disease based on their clinical status, and potentially prevent it through a data-driven selection of interventions[79]. Likewise, AI algorithms could recognize patients at a high risk of developing certain clinical disorders and send an alert to the patient's doctor. AI also has the potential to identify non-evident, clinically relevant patterns and advice for timely interventions. This can also contribute to a significant alleviation of side-effects, with less need for additional medications and less admissions to hospital. An overall change in clinical practice can also be induced by AI as these technologies can potentially identify patterns that distinguish fast-progressing patients from slow-progressing ones[80].

In addition, AI has the potential to transform recruitment in clinical trials, such as matching a patient to the most appropriate trial. Alternatively, high-quality data collected in a clinical trial environment are ideal to train AI systems to accurately predict outcomes. The use of "external" control arms exploiting real-world evidence (RWE) stored in vast clinical databases from routine clinical practice is of relevance to oncology, particularly in rare malignancies, where patient recruitment is quite difficult. In these cases, external control arms can be used as "*in silico*" comparators in early-phase clinical trials. RWE can also be of great value to analyse treatment efficiency in populations typically excluded from clinical trials (e.g., kidney or hepatic dysfunction, brain metastases, etc.) where the risk-benefit ratio is not known. Indication expansion is another area where RWE may be relevant, as it could provide a means of extracting data on clinical outcomes during off-label use, potentially broadening indications for approved therapies.

Healthcare Robotics

Although sometimes confused, AI and robotics are different from each other. Robotics is a technology that deals with the design and implementation of robots. The most popular concept of robots is that of machines programmed to perform a specific task over and over without any "intelligence". However, when combined with AI, the result is an artificially intelligent or smart robot that can be trained to carry out complex tasks that require more thought and adaptation. In smart robots, AI acts as the brain and the robotics acts as the body. Indeed, AI programs use data from the real world acquired through robotics to improve their performance; ergo, they help the robot strategize about future movements based on its prior motion. By doing so, the robot is learning how to interpret its own actions. Thus, automation, another subfield of AI, is concerned with how to train a robot to interact with the world around it in generalized and predictable ways, manipulate objects in interactive environments, and interact with people. Reinforced learning, which obviates the requirement of labelled data, is used to safely explore a policy space without committing errors that harm the system itself or others. Advances in reliable machine perception, including computer vision, force, and tactile sens-

ing, much of which is driven by machine learning, will continue to be key enablers to advance the capabilities of robotics. Indeed, robotics can potentially enhance human abilities in healthcare in surprising ways. "Surgical" robots are, for example, capable of performing complex surgical interventions, including minimally invasive and "surgeonless" surgical procedures. The Da Vinci system is the standard of care in multiple laparoscopic practices and is used in nearly three quarters of a million procedures a year[81,82]. Robotic surgery application strengthens and allows a higher productivity, accuracy, and efficiency of surgical procedures, favouring a faster recovery and improving patient outcomes.

Virtual Medicine

The human body is one of the biggest data platforms. Computational models can be used as virtual substitutes or representations of a subset of the patient's body or the whole (e.g. genetics, cell or tissue data, etc.)[83,84]. These virtual representations can integrate data from EHRs, telemedicine, wearables, and other platforms, tools, or media[85]. Available information should include past medical history, genetics, environment, diet, lifestyle, behaviours, preferences, socioeconomics, location, measurements from wearable technologies, and access and adherence to treatment, among others. MRI and CT scans can also be used to produce a virtual physiologic representation of the patient, reproducing the anatomy, tissue density, organ architecture and dynamic physiology of an individual patient. These virtual representations of a patient can then be used to develop a true personalized medicine for each individual patient[19,86-88]. In this sense, different systems are being created to represent clinical data in a reliable, ordered, and expandable formats, such as FHIR (Fast Healthcare Interoperability Resources), thereby simplifying data exchange across sites[89]. These virtual patients will therefore assist as a tool to support clinical decisions.

Limitations and Challenges

AI presents different challenges and limitations that must be addressed in the future. One of the major limitations lies in the fact that machine or deep learning systems can be considered as "black boxes" focused on predicting outputs from data, without taking into account the reasoning or justification by which a particular outcome is obtained[90]. Because of this, AI is not immune to the "garbage in, garbage out" problem. Undeniably, data and algorithms can be both flawed. Data can certainly be skewed because of the incompleteness and inaccuracy of existing datasets. Algorithms can be skewed, deliberately or not in ways that seek to achieve particular outcomes. Machines that learn on biased data will make biased decisions, posing new ethical problems and raising new questions. For example, how to prevent certain individuals or groups of people from using flawed automated systems to exploit them for their own benefit instead of maximizing their value for the public good? Ethics are the moral principles that govern an individual's behaviour or activity. The medical community must agree on appropriate ethical principles and frameworks to respond to these new challenges. Those frameworks are valuable tools as they set a basis that has moral authority to build global standards which reduce the risks that come with new technologies. Another weakness is the necessity of enormous datasets that may not be readily accessible[37]. As described above, AI frameworks built on inaccurately or insufficiently large training datasets can result in erroneous estimations and perpetuate biases. Thus, massively large datasets are essential for building accurate models. However, the majority of datasets are orders of magnitude too small for deep learning algorithms. It should also be stressed that AI is different from human intelligence in diverse ways; outperforming in one assignment does not necessarily suggest excelling in others. One of the biggest challenges is related to regulatory bodies, which have started to unlock the gate to machine-learning algorithms. To effectively apply AI platforms in the wider clinical care, regulatory bodies request that AI applications achieve results at least as well as experienced physicians. As the field matures, expectancies for AI accuracy will unsurprisingly become higher and higher.

Conclusions

Although highly successful in the past, the current medical R&D model is failing at improving health due to a series of flaws and defects inherent to the model itself. Therefore, we urgently need to generate a new collective intelligence, combining human and AI, to dramatically change the current model. Applying a global medical approach to the research flow, while incorporating the appropriate AI tools, will result in better disease prevention, earlier diagno-

sis and improved, cost-effective drug discovery, enhancing patient care standards through a truly personalized and holistic approach. Moreover, in an information-based environment, such as the current world, pharma and biotechnology companies, scientists and physicians will be able to share their skills and discoveries, contributing to the advancement of medical research. Even patients will have the chance of contributing to medical progress by voluntarily sharing their own anonymized medical records. Patients may stop being spectators and become main actors in the health arena. Thus, a "system-centric" model could become a "patient-centric model". Finally, thanks to the new wealth of information and knowledge that will become available through this new collective intelligence, health policies would change, emphasizing prevention over treatment and promoting the democratization of diagnosis, which will allow to detect health issues much earlier and therefore, treated at much lower costs with targeted drugs. The benefit could be greater in peripheral areas or in developing countries, without access to care centres of medium/high complexity, as primary care doctors will be able to make difficult diagnoses earlier and determine which patients should be referred to other centres to receive specialized care. In summary, this colossal common effort based on a new collective intelligence will exponentially improve the quality of medical research, resulting in a radical change for the better in the healthcare model. AI, without replacing us, is here to help us achieve the ambitious goal set by the WHO in the Alma Ata declaration of 1978: "Health for All"[91].

Conflict of Interest

Carlos M. Galmarini and Maximiliano Lucius are employees of Topazium Artificial Intelligence.

Authors' Contribution

Carlos M. Galmarini: conceptualization, writing, reviewing and editing draft preparation, approval of final version. Maximiliano Lucius: conceptualization, writing, reviewing and editing draft preparation, approval of final version.

References

1) WHO. Available at: www.who.int.
2) Galmarini CM. https://www.bbvaopenmind.com/en/technology/innovation/smarter-medical-research-through-new-collective-intelligence/.
3) Galmarini CM. Smarter Medical Research Through a New Collective Intelligence. Available at: https://www.bbvaopenmind.com/en/technology/artificial-intelligence/from-hippocrates-to-artificial-intelligence-moving-towards-a-collective-intelligence/.
4) Malone T. Superminds: the surprising power of people and computers thinking together. Hachette Book Group 2018.
5) Diprose W, Buist N. Artificial intelligence in medicine: humans need not apply? N Z Med J 2016; 129: 73-76.
6) Hosny A, Parmar C, Quackenbush J, Schwartz LH, Aerts H. Artificial intelligence in radiology. Nat Rev Cancer 2018; 18: 500-510.
7) Rodriguez MA, Kotagiri R, Buyya R. Detecting performance anomalies in scientific workflows using hierarchical temporal memory. Future Gener Comput Syst 2018; 88: 624-635.
8) Savage N. The measure of a man. Cell 2017; 169: 1159-1161.
9) Di Maio M, Basch E, Bryce J, Perrone F. Patient-reported outcomes in the evaluation of toxicity of anticancer treatments. Nat Rev Clin Oncol 2016; 13: 319-325.
10) FDA U. Guidance for industry: patient-reported outcome measures: use in medical product development to support labeling claims: draft guidance. Health Qual Life Outcomes 2006; 4: 79.
11) Howie L, Hirsch B, Locklear T, Abernethy AP. Assessing the value of patient-generated data to comparative effectiveness research. Health Aff (Millwood) 2014; 33: 1220-1228.
12) LeBlanc TW, Abernethy AP. Patient-reported outcomes in cancer care – hearing the patient voice at greater volume. Nat Rev Clin Oncol 2017; 14: 763-772.
13) Fung CH, Hays RD. Prospects and challenges in using patient-reported outcomes in clinical practice. Qual Life Res 2008; 17: 1297-1302.
14) Snyder CF, Aaronson NK, Choucair AK, Elliott TE, Greenhalgh J, Halyard MY, Hess R, Miller DM, Reeve BB, Santana M. Implementing patient-reported outcomes assessment in clinical practice: a review of the options and considerations. Qual Life Res 2012; 21: 1305-1314.
15) Basch E, Deal AM, Dueck AC, Scher HI, Kris MG, Hudis C, Schrag D. Overall Survival results of a trial assessing patient-reported outcomes for symptom monitoring during routine cancer treatment. JAMA 2017; 318: 197-198.
16) Camacho DM, Collins KM, Powers RK, Costello JC, Collins JJ. Next-generation machine learning for biological networks. Cell 2018; 173: 1581-1592.
17) Lohr S. For Big-Data Scientists, 'Janitor Work' Is Key Hurdle to Insights. Available at: https://www.nytimes.com/2014/08/18/technology/for-big-data-scientists-hurdle-to-insights-is-janitor-work.html.
18) Press G. https://www.forbes.com/sites/gilpress/2016/03/23/data-prepara-

tion-most-time-consuming-least-enjoyable-data-science-task-survey-says/.

19) Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, Liu PJ, Liu X, Marcus J, Sun M, Sundberg P, Yee H, Zhang K, Zhang Y, Flores G, Duggan GE, Irvine J, Le Q, Litsch K, Mossin A, Tansuwan J, Wang, Wexler J, Wilson J, Ludwig D, Volchenboum SL, Chou K, Pearson M, Madabushi S, Shah NH, Butte AJ, Howell MD, Cui C, Corrado GS, Dean J. Scalable and accurate deep learning with electronic health records. NPJ Digit Med 2018; 1: 18.

20) Goldstein BA, Navar AM, Pencina MJ, Ioannidis JP. Opportunities and challenges in developing risk prediction models with electronic health records data: a systematic review. J Am Med Inform Assoc 2017; 24: 198-208.

21) Drew BJ, Harris P, Zegre-Hemsey JK, Mammone T, Schindler D, Salas-Boni R, Bai Y, Tinoco A, Ding Q, Hu X. Insights into the problem of alarm fatigue with physiologic monitor devices: a comprehensive observational study of consecutive intensive care unit patients. PLoS One 2014; 9: e110274.

22) Maxmen A. AI researchers embrace Bitcoin technology to share medical data. Nature 2018; 555: 293-294.

23) Yu MK, Ma J, Fisher J, Kreisberg JF, Raphael BJ, Ideker T. Visible machine learning for biomedicine. Cell 2018; 173: 1562-1565.

24) Torkamani A, Andersen KG, Steinhubl SR, Topol EJ. High-definition medicine. Cell 2017; 170: 828-843.

25) Angermueller C, Parnamaa T, Parts L, Stegle O. Deep learning for computational biology. Mol Syst Biol 2016; 12: 878.

26) Hand DJ, Yu K. Idiot's Bayes - not so stupid after all? Int Stat Rev 2001; 69: 385-398

27) Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. Nature 2018; 559: 547-555.

28) Shakhnarovich G, Darrell T, Indyk P. Nearest-neighbor methods in learning and vision: theory and practice MIT Press, 2005.

29) Rokach L, Maimon O. Classification Trees. Springer US, 2010.

30) Shawe-Taylor J, Cristianini N. Kernel methods for pattern analysis cambridge univ. Press, 2004.

31) James G, Witten D, Hastie T, Tibshirani R. An introduction to statistical learning with applications in R. Springer US, 2013.

32) Rencher AC. Methods of multivariate analysis. John Wiley & Sons, INC, 2002.

33) Domingos P. A few useful things to know about machine learning. Commun ACM 2012; 55: 78-87.

34) LeCun Y, Bengio Y, Hinton G. Deep learning. Nature 2015; 521: 436-444.

35) Webb S. Deep learning for biology. Nature 2018; 554: 555-557.

36) Bengio Y, Delalleau O. On the expressive power of deep architectures. In: lecture notes in computer science. Springer Berlin Heidelberg, 2011.

37) Mamoshina P, Vieira A, Putin E, Zhavoronkov A. Applications of deep learning in biomedicine. Mol Pharm 2016; 13: 1445-1454.

38) Ciresan DC, Meier U, Gambardella LM, Schmidhuber J. Deep, big, simple neural nets for handwritten digit recognition. Neural Comput 2010; 22: 3207-3220.

39) Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science 2006; 313: 504-507.

40) Eickholt J, Cheng J. DNdisorder: predicting protein disorder using boosting and deep networks. BMC Bioinformatics 2013; 14: 88.

41) Spencer M, Eickholt J, Jianlin C. A deep learning network approach to ab initio protein secondary structure prediction. IEEE/ACM Trans Comput Biol Bioinform 2015; 12: 103-112.

42) Lipton ZC, Berkowitz J, Elkan C. A critical review of recurrent neural networks for sequence learning. arXiv 2015; 1506.00019.

43) Lipton ZC, Kale DC, Elkan C, Wetzell R. Learning to diagnose with LSTM recurrent neural networks. arXiv 2015; 1511.03677.

44) Guimaraes GL, Sanchez-Lengeling B, Outeiral C, Farias PL, AspuruGuzik A. Objective-reinforced generative adversarial networks (ORGAN) for sequence generation models. arXiv 2017; 1705.10843.

45) Hartnett K, Doré A. Best-Ever Algorithm Found for Huge Streams of Data. Available at: https://www.quantamagazine.org/best-ever-algorithm-found-for-huge-streams-of-data-20171024/.

46) Oprea TIT, A. Target, chemical and bioactivity databases - integration is key. Drug Discov Today Technol 2006; 3: 357-365.

47) Smalley E. AI-powered drug discovery captures pharma interest. Nat Biotechnol 2017; 35: 604-605.

48) Sterling T, Irwin JJ. ZINC 15--ligand discovery for everyone. J Chem Inf Model 2015; 55: 2324-2337.

49) Hill J, Mulholland G, Persson K, Seshadri R, Wolverton C, Meredig B. Materials science with large-scale data and informatics: unlocking new opportunities. MRS Bull 2016; 41: 399-409.

50) Pulido A, Chen L, Kaczorowski T, Holden D, Little MA, Chong SY, Slater BJ, McMahon DP, Bonillo B, Stackhouse CJ, Stephenson A, Kane CM, Clowes R, Hasell T, Cooper AI, Day GM. Functional materials discovery using energy-structure-function maps. Nature 2017; 543: 657-664.

51) Aliper A, Plis S, Artemov A, Ulloa A, Mamoshina P, Zhavoronkov A. Deep learning applications for predicting pharmacological properties of drugs and drug repurposing using transcriptomic data. Mol Pharm 2016; 13: 2524-2530.

52) Ragoza M, Hochuli J, Idrobo E, Sunseri J, Koes DR. Protein-ligand scoring with convolutional neural networks. J Chem Inf Model 2017; 57: 942-957.

53) Mayr A, Klambauer G, Unterthiner T, Steijaert M, Wegner JK, Ceulemans H, Clevert DA, Hochreiter S. Large-scale comparison of machine learning

methods for drug target prediction on ChEMBL. Chem Sci 2018; 9: 5441-5451.

54) Xu K, Ba J, Kiros R, Cho K, Courville A, Salakhutdinov R, Zemel R, Bengio Y. Show, attend and tell: neural image caption generation with visual attention. arXiv 2015; 1502.03044.

55) Van Noorden R. Software beats animal tests at predicting toxicity of chemicals. Nature 2018; 559: 163.

56) Azuaje F. Computational models for predicting drug responses in cancer research. Brief Bioinform 2017; 18: 820-829.

57) Costello JC, Heiser LM, Georgii E, Gonen M, Menden MP, Wang NJ, Bansal M, Ammad-ud-din M, Hintsanen P, Khan SA, Mpindi JP, Kallioniemi O, Honkela A, Aittokallio T, Wennerberg K, Collins JJ, Gallahan D, Singer D, Saez-Rodriguez J, Kaski S, Gray JW, Stolovitzky G. A community effort to assess and improve drug sensitivity prediction algorithms. Nat Biotechnol 2014; 32: 1202-1212.

58) Yip KY, Cheng C, Gerstein M. Machine learning and genome annotation: a match meant to be? Genome Biol 2013; 14: 205.

59) Dudley JT, Deshpande T, Butte AJ. Exploiting drug-disease relationships for computational drug repositioning. Brief Bioinform 2011; 12: 303-311.

60) Alipanahi B, Delong A, Weirauch MT, Frey BJ. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. Nat Biotechnol 2015; 33: 831-838.

61) Ballester PJ, Mitchell JB. A machine learning approach to predicting protein-ligand binding affinity with applications to molecular docking. Bioinformatics 2010; 26: 1169-1175.

62) Carro MS, Lim WK, Alvarez MJ, Bollo RJ, Zhao X, Snyder EY, Sulman EP, Anne SL, Doetsch F, Colman H, Lasorella A, Aldape K, Califano A, Iavarone A. The transcriptional network for mesenchymal transformation of brain tumours. Nature 2010; 463: 318-325.

63) Langille MG, Zaneveld J, Caporaso JG, McDonald D, Knights D, Reyes JA, Clemente JC, Burkepile DE, Vega Thurber RL, Knight R, Beiko RG, Huttenhower C. Predictive functional profiling of microbial communities using 16S rRNA marker gene sequences. Nat Biotechnol 2013; 31: 814-821.

64) Abe H, Macmahon H, Shiraishi J, Li Q, Engelmann R, Doi K. Computer-aided diagnosis in chest radiology. Semin Ultrasound CT MR 2004; 25: 432-437.

65) Korfiatis P, Kline TL, Coufalova L, Lachance DH, Parney IF, Carter RE, Buckner JC, Erickson BJ. MRI texture features as biomarkers to predict MGMT methylation status in glioblastomas. Med Phys 2016; 43: 2835-2844.

66) Wang J, Yang X, Cai H, Tan W, Jin C, Li L. Discrimination of breast cancer with microcalcifications on mammography by deep learning. Sci Rep 2017; 6: 27327.

67) Yao J, Dwyer A, Summers RM, Mollura DJ. Computer-aided diagnosis of pulmonary infections using texture analysis and support vector machine classification. Acad Radiol 2011; 18: 306-314.

68) Cruz-Roa A, Gilmore H, Basavanhally A, Feldman M, Ganesan S, Shih NNC, Tomaszewski J, Gonzalez FA, Madabhushi A. Accurate and reproducible invasive breast cancer detection in whole-slide images: A Deep Learning approach for quantifying tumor extent. Sci Rep 2017; 7: 46450.

69) Gareau DS, Correa da Rosa J, Yagerman S, Carucci JA, Gulati N, Hueto F, DeFazio JL, Suarez-Farinas M, Marghoob A, Krueger JG. Digital imaging biomarkers feed machine learning for melanoma screening. Exp Dermatol 2016; 26: 615-618.

70) Litjens G, Sanchez CI, Timofeeva N, Hermsen M, Nagtegaal I, Kovacs I, Hulsbergen-van de Kaa C, Bult P, van Ginneken B, van der Laak J. Deep learning as a tool for increased accuracy and efficiency of histopathological diagnosis. Sci Rep 2016; 6: 26286.

71) Ehteshami Bejnordi B, Veta M, Johannes van Diest P, van Ginneken B, Karssemeijer N, Litjens G, van der Laak J, Hermsen M, Manson QF, Balkenhol M, Geessink O, Stathonikos N, van Dijk MC, Bult P, Beca F, Beck AH, Wang D, Khosla A, Gargeya R, Irshad H, Zhong A, Dou Q, Li Q, Chen H, Lin HJ, Heng PA, Hass C, Bruni E, Wong Q, Halici U, Oner MU, Cetin-Atalay R, Berseth M, Khvatkov V, Vylegzhanin A, Kraus O, Shaban M, Rajpoot N, Awan R, Sirinukunwattana K, Qaiser T, Tsang YW, Tellez D, Annuscheit J, Hufnagl P, Valkonen M, Kartasalo K, Latonen L, Ruusuvuori P, Liimatainen K, Albarqouni S, Mungal B, George A, Demirci S, Navab N, Watanabe S, Seno S, Takenaka Y, Matsuda H, Ahmady Phoulady H, Kovalev V, Kalinovsky A, Liauchuk V, Bueno G, Fernandez-Carrobles MM, Serrano I, Deniz O, Racoceanu D, Venancio R. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. JAMA 2017; 318: 2199-2210.

72) Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017; 542: 115-118.

73) Leachman SA, Merlino G. Medicine: the final frontier in cancer diagnosis. Nature 2017; 542: 36-38.

74) Haenssle HA, Fink C, Schneiderbauer R, Toberer F, Buhl T, Blum A, Kalloo A, Hassen ABH, Thomas L, Enk A, Uhlmann L. Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. Ann Oncol 2018; 29: 1836-1842.

75) Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, Venugopalan S, Widner K, Madams T, Cuadros J, Kim R, Raman R, Nelson PC, Mega JL, Webster DR. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA 2016; 316: 2402-2410.

76) Kermany DS, Goldbaum M, Cai W, Valentim CCS, Liang H, Baxter SL, McKeown A, Yang G, Wu X, Yan F, Dong J, Prasadha MK, Pei J, Ting MYL, Zhu J, Li C, Hewett S, Ziyar I, Shi A, Zhang R, Zheng L, Hou R, Shi W, Fu X, Duan Y, Huu VAN, Wen C, Zhang ED, Zhang CL, Li O, Wang X, Singer MA, Sun X, Xu

J, Tafreshi A, Lewis MA, Xia H, Zhang K. Identifying medical diagnoses and treatable diseases by image-based deep learning. Cell 2018; 172: 1122-1131 e1129.

77) Maxmen A. Deep learning sharpens views of cells and genes. Nature 2018; 553: 9-10.

78) Poplin R, Varadarajan AV, Blumer K, Liu Y, McConnell MV, Corrado GS, Peng L, Webster DR. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. Nat Biomed Eng 2018; 2: 158-164.

79) Vemulapalli V, Qu J, Garren JM, Rodrigues LO, Kiebish MA, Sarangarajan R, Narain NR, Akmaev VR. Non-obvious correlations to disease management unraveled by Bayesian artificial intelligence analyses of CMS data. Artif Intell Med 2016; 74: 1-8.

80) Abelson S, Collord G, Ng SWK, Weissbrod O, Mendelson Cohen N, Niemeyer E, Barda N, Zuzarte PC, Heisler L, Sundaravadanam Y, Luben R, Hayat S, Wang TT, Zhao Z, Cirlan I, Pugh TJ, Soave D, Ng K, Latimer C, Hardy C, Raine K, Jones D, Hoult D, Britten A, McPherson JD, Johansson M, Mbabaali F, Eagles J, Miller JK, Pasternack D, Timms L, Krzyzanowski P, Awadalla P, Costa R, Segal E, Bratman SV, Beer P, Behjati S, Martincorena I, Wang JCY, Bowles KM, Quiros JR, Karakatsani A, La Vecchia C, Trichopoulou A, Salamanca-Fernandez E, Huerta JM, Barricarte A, Travis RC, Tumino R, Masala G, Boeing H, Panico S, Kaaks R, Kramer A, Sieri S, Riboli E, Vineis P, Foll M, McKay J, Polidoro S, Sala N, Khaw KT, Vermeulen R, Campbell PJ, Papaemmanuil E, Minden MD, Tanay A, Balicer RD, Wareham NJ, Gerstung M, Dick JE, Brennan P, Vassiliou GS, Shlush LI. Prediction of acute myeloid leukaemia risk in healthy individuals. Nature 2018; 559: 400-404.

81) Kim W, Choi Y, Cho JY, Yoon YS, Han HS. Robot single incision left lateral sectionectomy via da Vinci-(R) Xi Single Site & vaginal extraction of the specimen. Surg Oncol 2020; 44; 254-255.

82) Morelli L, Guadagni S, Caprili G, Di Candio G, Boggi U, Mosca F. Robotic right colectomy using the Da Vinci Single-Site(R) platform: case report. Int J Med Robot 2013; 9: 258-261.

83) Brown SA. Building SuperModels: emerging patient avatars for use in precision and systems medicine. Front Physiol 2015; 6: 318.

84) Brown SA. Principles for developing patient avatars in precision and systems medicine. Front Genet 2016; 6: 365.

85) Lippman H. How apps are changing family medicine. J Fam Pract 2013; 62: 362-367.

86) Highnam G, Mittelman D. Personal genomes and precision medicine. Genome Biol 2012; 13: 324.

87) Mirnezami R, Nicholson J, Darzi A. Preparing for precision medicine. N Engl J Med 2012; 366: 489-491.

88) Miotto R, Li L, Kidd BA, Dudley JT. Deep patient: an unsupervised representation to predict the future of patients from the electronic health records. Sci Rep 2016; 6: 26094.

89) Mandel JC, Kreda DA, Mandl KD, Kohane IS, Ramoni RB. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. J Am Med Inform Assoc 2016; 23: 899-908.

90) Kourou K, Exarchos TP, Exarchos KP, Karamouzis MV, Fotiadis DI. Machine learning applications in cancer prognosis and prediction. Comput Struct Biotechnol J 2014; 13: 8-17.

91) WHO. Declaration of Alma-Ata. Available at: https://www.who.int/publications/almaata_declaration_en.pdf.