

AI for Health

Examples from industry,
government, and academia





Table of Contents

Abstract	3
Introduction to artificial intelligence	4
Introduction to health data	7
Molecules: AI for understanding biology and developing therapies	10
Fundamental research	10
Drug development	11
Patients: AI for predicting disease and improving care	13
Diagnostics and outcome risks	13
Personalized medicine	14
Groups: AI for optimizing clinical trials and safeguarding public health	15
Clinical trials	15
Public health	17
Conclusion and outlook	18

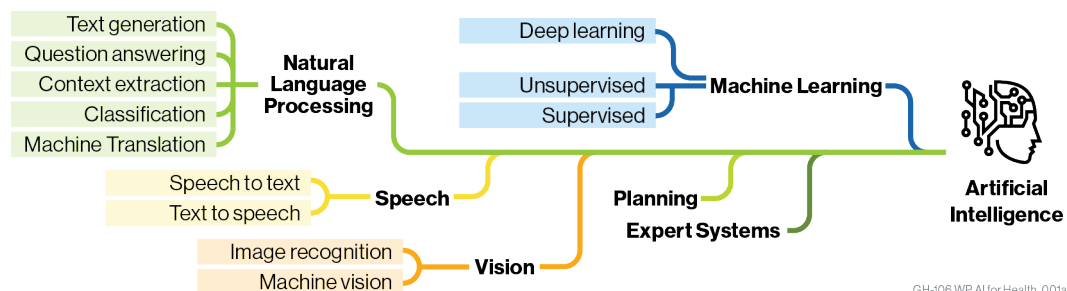
Abstract

This paper is directed toward a health-informed reader who is curious about the developments and potential of artificial intelligence (AI) in the health space, but could equally be read by AI practitioners curious about how their knowledge and methods are being used to advance human health. We present a brief, equation-free introduction to AI and its major subfields in order to provide a framework for understanding the technical context of the examples that follow. We discuss the various data sources available for questions of health and life sciences, as well as the unique challenges inherent to these fields. We then consider recent (past five years) applications of AI that have already had tangible, measurable impact to the advancement of biomedical knowledge and the development of new and improved treatments. These examples are organized by scale, ranging from the molecule (fundamental research and drug development) to the patient (diagnostics, risk-scoring, and personalized medicine) to the group (clinical trials and public health). Finally, we conclude with a brief summary and our outlook for the future of AI for health.

Introduction to artificial intelligence

Artificial intelligence (AI) is dramatically transforming industries across the board. At its core, AI combines elements from math and computer science to make sense of potentially massive and complicated datasets. AI techniques can be used to find patterns in seemingly unmanageable data, make predictions about the future given past outcomes, extract meaning from large volumes of text, and even digest pictures and sounds. We describe in this section some of the main theories and methods that underpin the use cases discussed later in the paper.

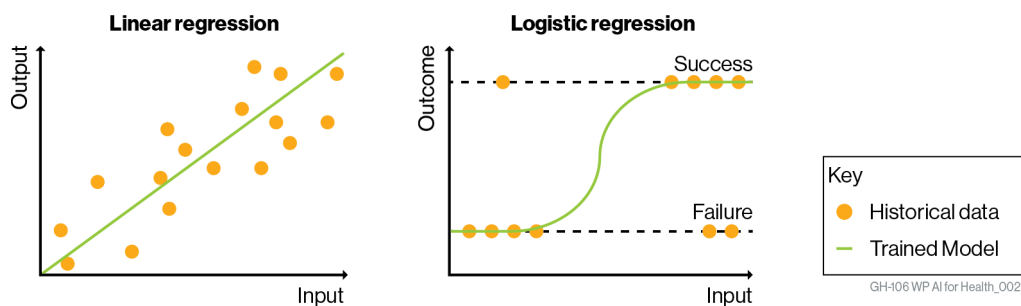
AI can be divided into a few subsets of techniques, the most mature of which include **machine learning**, **deep learning**, **computer vision**, and **natural language processing**. Although AI is a vast and rapidly evolving field, the theory behind it is unified through two mathematical concepts: **Bayesian statistics** and **optimization**. Bayesian statistics provide a framework for calculating probabilities by using observed data to tune theoretical statistical models. These methods therefore allow data scientists to connect real world data with abstract mathematical theories, and provide a high level of flexibility in the types and structures of data being modeled. **Bayesian networks**, for example, provide a way to analyze complex networks by representing probabilistic dependencies with graph theory. The other big mathematical idea behind AI is optimization, or finding the “best” set of conditions given an end goal. Often in AI, the desired result is to capture historical trends — in these cases, optimization methods guide algorithms to combine variables in ways that minimize the discrepancies between model predictions and past observations, thereby generating rules for making future predictions.




Machine learning (ML) is a subset of AI that encompasses supervised and unsupervised algorithms. **Unsupervised** models find patterns in large or messy data without needing “labels” or defined outcomes, while **supervised** models learn from input/output pairs to classify data into discrete groups or predict the outputs from a set of inputs.

One of the main applications of unsupervised modeling is **dimensionality reduction**, a set of techniques designed to transform datasets with many variables into modified representations with significantly fewer features. In some cases, these techniques can take hundreds of dimensions and recast the data to tens of features without sacrificing key insights present. Different algorithms achieve this, but one of the most popular is **principal component analysis**, which automatically extracts the most important combinations of individual variables to best capture the patterns in the data.

On the supervised side of ML, one of the most general algorithms is **linear regression**. In linear regression, the model adds input variables from historical data together in a way that best reproduces continuous outputs. When the model is presented with new data, it makes the most reasonable prediction using the combination rules learned from historical information. Like linear regression, logistic regression also finds the best combination of variables to predict an outcome, but **logistic regressions** predict binary outcome probabilities, such as the probability of a success versus a failure. They accomplish this by transforming input variables such that the output lies between 0 (negative outcome) and 1 (positive outcome).



Recent advances in computing power have allowed for the development of larger, more complex models that together make up the subset of AI known as **deep learning** (DL). DL is largely focused on **neural networks**, which are models built from series of computational “cells” designed to operate like a brain. Each individual cell processes its input in a different way, and groups of cells called “layers” work together to store and digest data. These layers then link to subsequent layers and transmit information through different firing mechanisms, much like how neurons transmit electrochemical signals throughout the nervous system. The end effect is that a network learns to associate inputs with outputs. The more layers involved in the network, the “deeper” the model.



Neural networks are especially good at learning to store raw nonnumerical data as computer-readable formats. Recent advances in DL now allow image and video data to be processed through **convolutional neural networks** (CNNs). These networks identify and encode relevant pixel features to create numeric representations of images; these advances have led to the emergence of **computer vision** as its own subfield of AI. A related set of neural networks are called **recurrent neural networks** (RNNs); these are well-suited to time-series and text data. In these models, basic neural network cells are placed in loops that combine subparts of the raw data. RNNs exploit the order in the source data (like word order in text or peaks and valleys in signals) to extract the meaning of the underlying text or sequence. CNNs and RNNs can both create meaningful encodings or representations of unstructured, complex data; these representations can be fed to other models to ask regression or classification questions, ultimately tying the raw sources to useful insight.

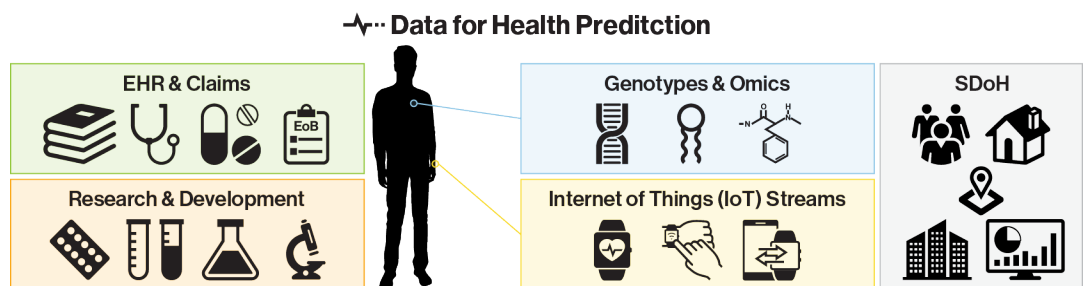
A final rapidly developing domain of AI is **natural language processing** (NLP). This family of methods allows data scientists both to numerically encode and extract insights from raw, unstructured text data (for example clinical notes or internet searches). These techniques can be used on their own to find useful patterns in, or classify large amounts of, text. They can also be used upstream of a machine-learning model by generating features from text data, such as keyword extraction or sentiment analysis. NLP methods often leverage neural networks behind the scenes, as they require transformation of nonnumeric sequential data into representations that a computer can use.

Introduction to health data

One of the areas enjoying some of the highest degree of AI innovation is health, taken here to mean both life sciences and healthcare delivery. Health is a huge market and is attractive to AI researchers given the tremendous volume and variety of health data being produced constantly, as well as the potential for improving care. At the same time, the health industry possesses unique challenges to overcome before meaningful progress can be achieved.

The scale of data surrounding people's health is constantly increasing, both in depth and in breadth. AI algorithms today can leverage a wide variety of sources, including:

- ▶ **Patient records:** Health encounters, ranging from routine immunizations to emergency surgeries, are now digitized and stored in Electronic Health Records (EHRs).
- ▶ **Administrative records:** Massive sets of claims from providers to payers tie together patient characteristics, diagnoses, prescriptions, and procedures.
- ▶ **Social determinants of health (SDoH):** Sources such as the Census Bureau offer a complementary lens into patient health journeys by highlighting demographic and socioeconomic factors.
- ▶ **Internet of Things (IoT) streams:** Wearables and smartphones provide constant and real-time signals via sensors that track vital signs such as heart rate, activity level, and blood oxygen saturation.
- ▶ **Genotypes:** Entire genetic maps, and hence knowledge of molecular predispositions, of patients are now routinely available due to the markedly reduced cost of sequencing the human genome (down from roughly \$100 million in 2001 to less than \$1,000 in 2019).¹
- ▶ **Omics:** Large-scale molecular fingerprints at the cellular and person levels, including, for example, libraries of proteins (proteomics), metabolites (metabolomics), and lipids (lipidomics), are made possible by the lower costs of data and processing power, coupled with advances in biochemistry and related sciences.
- ▶ **Research and Development (R&D):** Academic and commercial entities drive biological innovation and breakthroughs every day in the areas of fundamental biology, drug development, and translational applications, and in doing so generate large amounts of scientific data.



¹The Cost of Sequencing a Human Genome. National Human Genome Research Institute. <https://www.genome.gov/about-genomics/fact-sheets/Sequencing-Human-Genome-cost> (2019).

This multidimensional slew of data, coupled with modern developments in statistical theory, programming, and computational hardware, means that it is a truly exciting time to be applying AI to questions of health. That said, it is important to realize that there are some very real challenges to accessing the data required for, or developing predictive models about, biomedical problems, including:

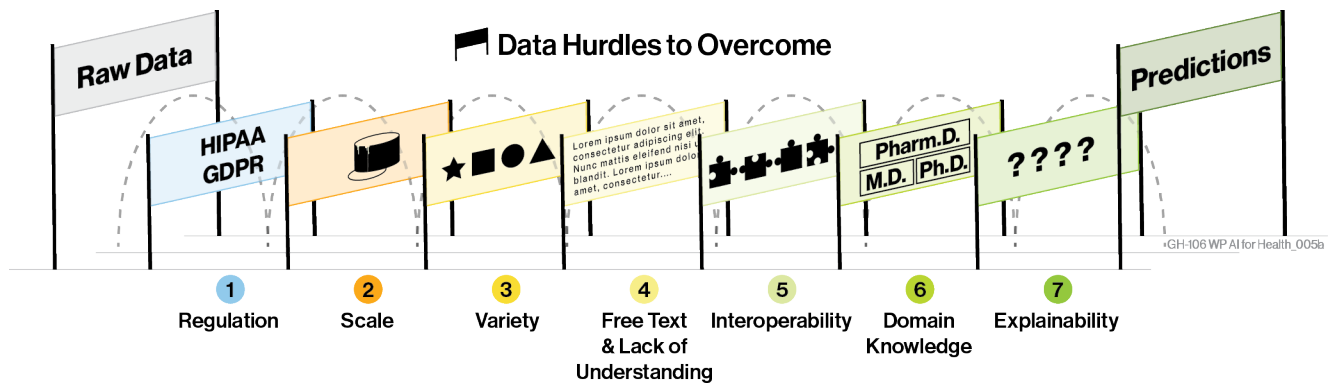
- ▶ **Regulatory considerations:** Protected health information (PHI) is strictly regulated by the Health Insurance Portability and Accountability Act in the United States and the General Data Protection Regulation in Europe, including criminal penalties and hefty fines for violations, such that sharing patient data involves significant logistical overhead.^{2,3}
- ▶ **Scale of data:** Health data are huge, and growing faster than any other industry;⁴ the sizes of these data pose expensive barriers to storage and analysis, as the data available to query and train a model can outstrip conventional computational memory.
- ▶ **Variety of data:** The huge diversity of data sources and file types means that it is difficult to develop universal models for storing data in structured and connected databases.
- ▶ **Free text and lack of standardized ontologies:** Often the most medically interesting information lies in free text (for example, doctor's notes or clinical trial inclusion/exclusion criteria) which, when coupled with the lack of globally approved, freely available ontologies for various concepts such as procedures or medical histories, renders it challenging to transform raw data into coded structure that AI algorithms can exploit.
- ▶ **Interoperability challenges:** The sheer number of competing players in the health market poses a blocker to universal interoperability of datasets: For example, it was found that in 2018, health systems had on average 18 different EHR vendors.⁵
- ▶ **Domain knowledge:** Health data often refer to concepts that require highly specialized (i.e., professional or doctoral) knowledge to fully understand and build useful model features from.
- ▶ **Need for white-box models:** To impact patient health, the main contributors to a prediction must be known, such that explainable ("white box") models are desired even more so than in other industries; furthermore, healthcare professionals are trained in evidence-based thinking and hence are unlikely to trust or use unexplainable predictions.

²European Parliament. Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation – GDPR). International and European Labour Law 958–981 (2016). doi:10.5771/9783845266190-974.

³Summary of the HIPAA Privacy Rule. U.S. Department of Health and Human Services. <http://www.hhs.gov/ocr/privacy/hipaa/understanding/summary/index.html> (2020).

⁴Reinsel, D., Gantz, J. & Rydning, J. Data Age 2025: The Digitization of the World From Edge to Core. Int. Data Corp. 28 (2018).

⁵Sullivan, T. Why EHR data interoperability is such a mess in 3 charts. Healthcare IT News. <https://www.healthcareitnews.com/news/why-ehr-data-interoperability-such-mess-3-charts> (2018).




The above considerations are by no means meant to be exhaustive of the data sources and challenges in the health space, but rather are included to introduce and highlight main themes. The remainder of this paper offers a sample of recent use cases in which AI has had measurable impact, organized in order of increasing scale from the molecule to the community.

Molecules: AI for understanding biology and developing therapies

Applied to the nanoscopic and microscopic worlds, AI can help scientists push our understanding of biology and generate novel therapies against diseases. At this scale, algorithms can make use of the massive amounts of data generated by researchers to learn mechanisms and propose new hypotheses and potential treatments.

Fundamental research

One area in which AI, particularly deep learning, has made its mark is protein folding. Proteins are the components of cells that allow organisms to function and are responsible for an astoundingly diverse set of biological tasks, from immune responses to regulation of cellular division. Molecularly, proteins are strings of building blocks called amino acids, and these strings must fold into specific three-dimensional shapes, called native states, before they are biologically active. The thermodynamic hypothesis broadly states that a protein's specific sequence of amino acids determines its native state; in other words, one can predict a protein's structure just by knowing its sequence.⁶ This is very medically relevant given that specific sites on proteins are often the target of drugs, and that many debilitating diseases of aging such as Alzheimer's or Parkinson's are linked to improperly shaped proteins.⁷



In practice, developing scientific methods to predict protein folding has proven an immense challenge. While significant progress has been made in the past few decades, state-of-the-art methods still cannot fully predict shape from sequence. One of the most promising recent developments, however, lies in the use of deep learning and neural networks to tackle this problem. In 2018, Google DeepMind uncovered AlphaFold, a DL pipeline that uses several different neural networks to predict the structure of a protein given a sequence input. It does this in a series of steps that includes comparing the test sequence to a library of known sequences and structures, guessing reasonable conformations for subsets of the input protein, combining the guesses together in different ways, and then figuring out what the best-guess shape is. The accuracy of the results roughly doubled the trend of progress in the field.^{8,9} In this case, AI helped researchers get closer to achieving their goal of understanding a protein's shape based on its structure, which would be of tremendous value for rational drug design in which molecules are generated to target specific geometries of specific molecular entities.

⁶Govindarajan, S. & Goldstein, R. A. On the thermodynamic hypothesis of protein folding. *Proc. Natl. Acad. Sci.* 95, 5545–5549 (1998).

⁷Dobson, C. M. Protein folding and misfolding. *Nature* 426, 884–890 (2003).

⁸AlQuraishi, M. AlphaFold at CASP13. *Bioinformatics* 35, 4862–4865 (2019).

⁹Evans, R. et al. De novo structure prediction with deep-learning based scoring. Thirteen. *Crit. Assess. Tech. Protein Struct. Predict. Abstr.* 1–4 (2018).

Drug development

A particularly active area for AI within health is the drug development process, with the number of funded startups active (currently at about 186) in this sphere growing about 400% between November 2017 and January 2020.¹⁰ The opportunities for application of AI span the entire development pipeline, from fundamental chemical reactions to experimental design to identification of new therapies.



In 2019, scientists at Pfizer and researchers at the University of Cambridge developed a neural network that could make synthesis predictions for pharmacologically relevant compounds. Their tool works both for the problems of reaction prediction (forward from reactants to products) and retrosynthetic analysis (backward from products to reactants), and their method beat professional human chemists by achieving an accuracy of about 90% for well-represented reactions (versus about 80% human accuracy). Much like the AlphaFold network learns protein folds from protein sequences, this reaction network learns the properties of specific reactants that lead to specific products, and vice versa, leveraging the large amount of knowledge generated by organic chemists in the past few centuries.^{11,12} By automating the design of chemical reactions, AI could help pharmaceutical companies design drug syntheses and optimize drug production processes.

On a more applied level, in 2018 a team of researchers from the Universities of Cambridge, Manchester, and Aberystwyth discovered novel antimalarial properties of triclosan, an antibacterial compound found in products such as toothpaste. This was achieved using an AI-powered robot named Eve that automated the scientific process by using ML to generate hypotheses, run automated experiments in yeast cells, and analyze the results. Behind the scenes, Eve uses data from positive experiments (hits) to run linear regressions and learn quantitative structure-activity relationships that quantify a compound's activity in a specific assay, given its chemical properties such as size or hydrophobicity. Triclosan works by inhibiting two different malarial enzymes, including one that is the target of the current standard-of-care drug

pyrimethamine. These two different mechanisms of action render triclosan a "polypharmacologic" compound since it can target the malaria parasite at both the blood and liver stages of its life cycle. There is a growing need for such new therapeutic discoveries as drug resistance against standard medicines, like pyrimethamine, grows in endemic regions such as Africa.¹³⁻¹⁵

As a final example, BERG, a clinical-stage biotechnology startup near Boston, identified a candidate drug for advanced pancreatic cancer that in November 2017 was awarded orphan-drug designation by the Food and Drug Administration (FDA) and in summer 2019 completed a Phase 2 trial. This drug works by undoing metabolic changes associated with cancer at the cellular level to restore normal energy pathways, thereby turning normal protective mechanisms such as apoptosis (controlled, prophylactic cell death) back on. BERG has developed a technology that combines huge amounts of patient sample data (including genomics, proteomics, metabolomics and lipidomics) with medical data such as EHR information to learn patient-level biological graphs using a technique called Bayesian networks. From these graph representations of patients, the team can identify disease biomarkers and critical molecules against which they can virtually screen therapies for toxicity and efficacy by seeing where the biological networks tend to bunch up or by seeing how diseased networks differ from healthy networks. Through their AI-based pipeline, then, they can reverse the hypothesis generation framework by going from patient-level results to biological mechanisms that, because of their empirically derived natures, are likely to succeed in the clinic.¹⁶⁻¹⁸

¹⁰ Smith, Simon. 186 Startups Using Artificial Intelligence in Drug Discovery. *BenchSci*. <https://blog.benchsci.com/startups-using-artificial-intelligence-in-drug-discovery> (2020).

¹¹ Schwaller, P. et al. Molecular Transformer: A Model for Uncertainty-Calibrated Chemical Reaction Prediction. *ACS Cent. Sci.* 5, 1572–1583 (2019).

¹² Lee, Alpha A. et al. Molecular transformer unifies reaction prediction and retrosynthesis across pharma chemical space. *Chem. Commun.* 55, 12152–12155 (2019).

¹³ Bilsland, E. et al. Yeast-based automated high-throughput screens to identify anti-parasitic lead compounds. *Open Biol.* 3, 120158 (2013).

¹⁴ Bilsland, E. et al. Plasmodium dihydrofolate reductase is a second enzyme target for the antimalarial action of triclosan. *Sci. Rep.* 8, 1038 (2013).

¹⁵ AI 'scientist' finds that toothpaste ingredient may help fight drug-resistant malaria. University of Cambridge Research News. <https://www.cam.ac.uk/research/news/ai-scientist-finds-that-toothpaste-ingredient-may-help-fight-drug-resistant-malaria> (2018).


¹⁶ Zhang, L., Rodrigues, L. O., Narain, N. R. & Akmaev, V. R. bAlcis: A Novel Bayesian Network Structural Learning Algorithm and Its Comprehensive Performance Evaluation Against Open-Source Software. *J. Comput. Biol.* 26, 1–11 (2019).

¹⁷ Our Approach. BERG <https://www.berghealth.com/research/healthcare-professionals> (2020).

¹⁸ Hendifar, A. E. et al. Phase I study of BPM 31510 in advanced solid tumors: Updated analysis of a novel treatment with promising activity. *J. Clin. Oncol.* 30, 3015–3015 (2012).

Patients: AI for predicting disease and improving care

Moving up to the person scale, AI is increasingly being applied to the patient experience, including risk prediction, diagnosis, and treatment. New algorithms exploit the recent explosion of patient health data availability to generate predictions that complement services from trained providers.



Diagnostics and outcome risks

In April 2018, the software system IDx-DR became the first fully autonomous AI system to gain FDA approval. IDx-DR reads patient eye scans and calculates the probability that the scans show evidence of diabetic retinopathy, a disease of the small blood vessels in the eye due to diabetes and the leading cause of blindness for working-age Americans. The software triages patients into a high-risk group, who are then referred to ophthalmic specialists, and a low-risk group, who are scheduled for rescreening in a year. Much like an ophthalmologist, the algorithm functions by learning how to read pictures of retinas and what features of the scans indicate disease by using a convolutional neural network. With 30 million Americans suffering from diabetes, and about 30% of them affected by diabetic retinopathy, being able to accurately score patients on risk results in greater efficiency of specialty care for only those who require it.^{19,20}

Another application of AI-based diagnosis is from integration with IoT devices, namely the Apple Watch. Researchers at the University of California, San Francisco in 2018 conducted experiments to see if this wearable coupled with a deep learning app could diagnose atrial fibrillation (AF). AF is an anomalous heart rhythm and is the leading cause of stroke. Traditional diagnosis involves a 12-lead electrocardiogram read by a specialist, but the study authors showed that using heart rate and step data from the watch sensor could train a recurrent neural network to be able to identify AF episodes. Here the software teaches itself to read heart scans from lots of patient scans, finding both normal cardiac patterns and anomalies associated with AF.²¹ By moving the point of care from the clinic to wherever the patient is, AI can deliver the benefits of constant monitoring of at-risk individuals without incurring the costs associated with traveling to, and being seen by, scarce medical providers.

¹⁹ FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems. U.S. Food and Drug Administration. <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye>.

²⁰ Gulshan, V. et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA 316, 2402 (2016).

²¹ Tison, G. H. et al. Passive Detection of Atrial Fibrillation Using a Commercially Available Smartwatch. JAMA Cardiol. 3, 409 (2018).

AI is also being used to calculate the risk of patients developing certain conditions. One example comes from the field of cardiology: scientists at the University of Nottingham in 2017 were able to use EHR data and demographic information on patients to predict the risk of cardiovascular events. They tried several of the standard ML techniques, including logistic regression and neural networks, with the latter being the most successful, and all beating the standard algorithm used by the American College of Cardiology.²² While preliminary, these results demonstrate that the information present in clinical notes, combined with social determinants of health, can teach relatively simple models to be as predictive as trained experts in triaging patients. Freeing up specialists' time by performing preliminary triaging, scaled up by a doctor's entire patient load, can have real measurable impact on provider economics and efficiency by increasing the time, and hence quality of care, the provider is able to deliver to their most at-risk patients.

Personalized medicine

Developments in genomics and medicine in the past century have highlighted the fact that many disorders, including rare and genetic diseases, are not responsive to traditional one-size-fits-all therapies. In other words, the same drug on two different patients presenting with similar symptomologies can have very different effects. Furthermore, quickly mutating diseases such as aggressive cancers or antibacterial-resistant infections require new interventions as the disease evolves. To meet such needs, personalized (or precision) medicine has evolved to provide targeted therapies that work for specific patient situations. The fundamental challenge of personalized medicine is to predict or develop the most effective intervention, given data (medical history, diagnostics, omics, etc.) about a patient and their disease. As such, AI finds natural applications here.

Alcoholism poses a major burden, affecting about 5% of the global population, yet its psychiatric etiology is complex and not completely understood. There are currently only a handful of medicines that are approved to treat alcohol

dependence, but none are extremely effective at a large scale. As an application of personalized medicine, researchers at the Mayo Clinic performed a study in 2017 in which they used ML to tie patient metabolomics to their likelihood to respond to acamprosate, one of the drugs approved for treatment of alcoholism. They measured patient demographics, metabolite levels, and baseline alcohol cravings before treatment, and then correlated these variables to the probability of response to treatment by using logistic regression. The study was able to generate predictive models with reasonable predictive power and showed that levels of the amino acid aspartate was a consistently good predictor of treatment effectiveness.²³ This study demonstrates that AI can generate useful personalized predictions about patients' individual responses to interventions, thereby increasing quality of care by avoiding spending time and money on therapies unlikely to be effective.

In 2011, the U.S. Department of Veterans Affairs (VA) launched the Million Veteran Program, an effort to collect genomic information on a million veterans so that these data can be correlated with information on military service, lifestyle, and medical care to better understand the interplay between genotype, environment, and phenotype. As the United States' largest integrated health system and the world's largest collector of genomic data linked to health data, the VA has offered the AI community a tremendous opportunity, and recently efforts have started to pay off. In 2018, researchers from around the world representing institutions such as Harvard, Stanford, and Cambridge Universities worked with VA scientists to help link genomic data on hundreds of thousands of veterans to their lipid levels, gleaned key insights into genetic underpinnings of cardiovascular and metabolic diseases. Several ML techniques were used in this work, including principal component analysis (PCA) to deal with the vast amounts of data present in a patient's genome. By representing patient ancestries as PCA-reduced variables of genomes, combined with lab results on lipid levels, EHR records of

²² Weng, S. F., Reips, J., Kai, J., Garibaldi, J. M. & Qureshi, N. Can machine-learning improve cardiovascular risk prediction using routine clinical data? PLOS ONE 12, e0174944 (2017).

²³ Hinton, D. J. et al. Metabolomics biomarkers to predict acamprosate treatment response in alcohol-dependent subjects. Sci. Rep. 7, 2496 (2017).


cardiovascular disease, and demographic data, the team was able to use logistic regression to extract the effects of specific genetic mutations on coronary artery disease (CAD). They identified that the inactivation of one gene, PDE3B, was associated with lower risk of CAD. The effects of inactivating this gene are strikingly similar to the molecular mechanism of action of cilostazol, a drug that is approved by the FDA for the treatment of insufficient blood flow in the legs, but not for CAD. By making this connection, uncovered through ML techniques on genomic and health data, the researchers then suggested that cilostazol may be a therapeutic possibility against CAD,^{24,25} and recent clinical studies have emerged with results supporting this hypothesis.^{26,27} This case study shows the potential of applying AI to linked genomic and health data to understand at an individual level the genetic driving forces behind disease, and further to generate new drug candidates from existing medicines.

Groups: AI for optimizing clinical trials and safeguarding public health

The last level of scale we consider here is the group, or how AI can support cohort and population outcomes. Two areas of development include clinical trials and public health; in each of these contexts, AI models are being used to make sense of complicated, disparate sets of data about individuals to make predictions that can benefit larger groups.

Clinical trials

Clinical trials are an enormous business operation: In 2013, Phase 3 trials in the United States alone cost around \$10 billion. The FDA approval process is strict and rigorous for good reason, but this means that each new therapy costs about \$1.5 billion-\$2 billion to bring to market when one accounts for R&D costs and the costs of failed trials. In the end, only 10% of candidate therapies that make it to Phase 1 make it to market, and one of the major factors for trial failure is the inability to recruit enough eligible patients who do not attrite before trial completion.^{28,29}



²⁴ Million Veteran Program (MVP). U.S. Department of Veterans Affairs. <https://www.research.va.gov/mvp/> (2017).

²⁵ Klarin, D. et al. Genetics of blood lipids among ~300,000 multi-ethnic participants of the Million Veteran Program. *Nat. Genet.* 50, 1514–1523 (2018).

²⁶ Lee, D.-H. et al. Effect of cilostazol, a phosphodiesterase-3 inhibitor, on coronary artery stenosis and plaque characteristics in patients with type 2 diabetes: ESCAPE study. *Diabetes Obes. Metab.* 21, 1409–1418 (2019).

²⁷ Chao, T.-H., Tseng, S.-Y., Liu, P.-Y. & Li, Y.-H. P1929 A randomized controlled trial evaluating outcome impact of cilostazol in patients with coronary artery disease and at a high risk of cardiovascular disease. *Eur. Heart J.* 40, ehz748–0676 (2019).

²⁸ Harrer, S., Shah, P., Antony, B. & Hu, J. Artificial Intelligence for Clinical Trial Design. *Trends Pharmacol. Sci.* 40, 577–591 (2019).

²⁹ Fogel, D. B. Factors associated with clinical trials that fail and opportunities for improving the likelihood of success: A review. *Contemp. Clin. Trials Commun.* 11, 156–164 (2018).

The United States has a central repository of all clinical trials, clinicaltrials.gov, but this resource doesn't translate directly into streamlined enrollment, in part because the inclusion/exclusion trial criteria are in free text, use medical jargon, and do not follow a standardized ontology. To address these problems, researchers at Janssen Pharmaceuticals and Columbia University used natural language processing (NLP) techniques to develop the Criteria2Query algorithm, which takes as input a free-text description of clinical trial eligibility and outputs a standardized collection of entities (condition, drug, measurement, procedure, or observation), attributes (values and time occurrences) and negation statuses (whether presence of the entity/attribute pair determines eligibility or ineligibility). For example, inputting "13-15 years old" would yield an eligibility entity of "age" with value attributes "13-15 years." The algorithm also decodes medical abbreviations and can go from a free-text description to a standard ontological entry. For example, "AD" maps to "Alzheimer's Disease" and then matches the International Classification of Diseases-Tenth Revision (ICD-10) code G30. Criteria2Query does this through an intricate set of NLP steps, including paragraph and sentence parsing, named entity recognition, negation detection, relation extraction, logic detection, and attribute normalization.³⁰ By using AI to translate unstructured criteria defining each of the many clinical trials into well-formatted and standardized outputs, tools can be developed to help patients and providers find experimental treatments through large-scale querying of available trials.

Besides simply cleaning and understanding eligibility criteria, a major challenge of trial design is the actual matching of eligible patients to relevant trials. Deep 6 AI is a California-based startup working to address this need. Its team accomplishes this by using NLP and ML to sift through data buried in patient health records, analyzing both structured (diagnosis codes, procedure codes, etc.) and unstructured (doctor notes, pathology reports, etc.) information to reconcile disparate sources of health indicators into one multidimensional representation of the clinical state of a patient. Having a collected

view of a patient then allows them to identify possibilities for intervention, including for conditions not present explicitly in patient data. Their algorithms match these opportunities to currently available trials, thereby greatly speeding up the patient recruitment process; in one hour they were able to find 16 subjects for a trial, whereas a typical recruitment process had taken six months to find two patients.^{31,32}

A last example in the context of clinical trials is the application of AI to design of the interventions themselves. Clinical trial design involves many simultaneously moving parts, such as defining appropriate subject inclusion/exclusion criteria or trial endpoints. These choices are often delicate balancing acts: for example, if a trial's eligibility is too strict, not enough patients will be recruited and hence intervention significance will be difficult to attain because of low statistical power. On the other hand, if the eligibility is too lenient, the trial will be flooded by subjects who may not respond as predicted and significance will be difficult to attain due to seemingly low effect size, in addition to increased trial costs. As a result, an efficacious drug may appear not to be efficacious. To address this, a startup called [trials.ai](https://www.trials.ai) in San Diego is using ML and NLP techniques to optimize the design of the trial protocols. Their product mines data such as published journal articles and proprietary pharmaceutical or medtech client data to, for example, suggest appropriate thresholds for patient recruitment criteria.^{31,33} By being able to fine-tune clinical trials to optimize economics and design for the most relevant patient populations, AI has the potential to help new drugs get to market faster and cheaper.

³⁰ Yuan, C. et al. Criteria2Query: a natural language interface to clinical databases for cohort definition. *J. Am. Med. Inform. Assoc.* 26, 294–305 (2019).

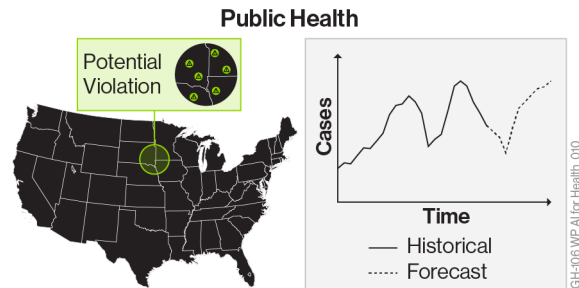
³¹ Woo, M. Trial by artificial intelligence. *Nature* 573, S100–S102 (2019).

³² How It Works. Deep 6 AI <https://deep6.ai/how-it-works> (2020).

³³ How We Do It. [Trials.ai https://www.trials.ai/](https://www.trials.ai/) (2020).

Public health

The potential use cases for AI in health do not stop at the cohort level — they can also be applied to much larger groups such as city, or even national, populations to assist public health efforts.



Google teamed up with researchers from Harvard and the public health departments of Chicago and Las Vegas between 2016-2017 to predict sources of food-borne illnesses and identify health violations in restaurants from Google search queries. To do so, the team created a two-step model. In the first step, a model uses NLP to encode the search strings followed by supervised ML classification to predict whether the query was about food-borne illness. In the second step, the model uses location data from opted-in users to estimate which restaurants the user had visited. Combining the web search with the location model, and aggregating over all the unique users, allowed for a prediction of which restaurants were correlated with searches likely to indicate a food-borne pathogen. Leveraging these predictions, both public health departments increased their inspection efficiency by more than a factor of 2 over their typical routine site inspection processes.³⁴ Public health departments are often underfunded and understaffed, so unlocking such efficiency gains with AI can provide measurable benefits to the community by freeing up public health officials to focus their time on other high-impact activities.

AI methods have also found application in epidemic surveillance and forecasting. Through their FluSight challenge, the U.S. Centers for Disease Control and Prevention (CDC) has teamed up with academic and commercial partners to use influenza incidence data collected at reporting sites across the country to forecast the flu season. The winning team in the 2018 season was a group from the U.S. Department of Energy's Los Alamos National Laboratory that used Bayesian statistical models of flu infection time series to learn trends at the various geographical levels. Importantly, their model was able to beat the other competitors, some of whose models used additional data sources, without incorporating orthogonal information such as weather patterns or web traffic. This demonstrates that historical information, combined with innovative ML, is sufficient to generate actionable public health predictions and provide huge value to constituents. The ability to forecast influenza spread can allow the CDC to optimally time vaccination drives, prepare antiviral supply chains ahead of peak clinical need, and dynamically allocate resources to regions of the country most at risk for epidemic.^{35,36}

³⁴ Sadilek, A. et al. Machine-learned epidemiology: real-time detection of foodborne illness at scale. *Npj Digit. Med.* 1, 36 (2018).

³⁵ FluSight: Flu Forecasting. U.S. Centers for Disease Control and Prevention <https://www.cdc.gov/flu/weekly/flu-sight/index.html>.

³⁶ Osthus, D. & Moran, K. R. Multiscale Influenza Forecasting. *ArXiv: 1909.13766 [stat.AP]* (2019).

Conclusion and outlook

At its core, AI is a family of methods that find patterns in, and make predictions about, data. Advances in the biomedical sciences have vastly increased the scale and variety of data available in the health and life sciences sectors, while parallel advances in data science and computing have created new algorithms to process, analyze, and extract measurable insight from various sources of information. While there are very real challenges to overcome when working with health data, the potential for AI to advance medical knowledge and care is vast and present at all scales, from the molecule to the community. By leveraging AI, providers can triage patients to focus their time on more difficult cases, life scientists can discover new medicines, patients can find relevant clinical trials, and public health agencies can more proactively serve their constituents. We have provided here examples of all these applications of AI from the past five years, and we anticipate that opportunities for AI to improve health will continue to grow and develop in the near future.


While there are many exciting developments underway and on the horizon for AI in health, it is important to realize that AI is not a panacea. Namely the following limitations are especially important to keep in mind:

- ▶ **Biology will not be fully solved by models.** Models can only find patterns in already observed measurements or correlate historical inputs with outputs, so they will never displace the critical role that researchers have in generating and testing novel scientific hypotheses.
- ▶ **Clinician expertise and experience will not be supplanted by robot doctors.** Medical care demands empathy and a mastery of understanding all aspects of the situational nuances of each patient; no model can ever replicate these skills that experience in patient care alone can provide.
- ▶ **Patient health trajectories will not be fully and perfectly predictable.** An individual's environment and their personal choices are a priori unknowable but carry tremendous weight in affecting their health, implying that computational algorithms will never be able to produce deterministic, exact forecasts.
- ▶ **AI models will not be fully explainable.** Often the most sophisticated, powerful AI models suffer from a lack of transparency in how exactly they use data to make predictions; in health this is a salient shortcoming, as it is difficult to intervene if the root causes of a disease state are not known.
- ▶ **AI technology will not be free of bias or equitably accessible.** The output of an AI model always reflects the input data, so model forecasts necessarily reflect the biases inherent to the data being collected; put another way, until medical care is equally accessible to all, AI tools will only accurately predict outcomes for those populations able to access care and generate training data in the first place.

Despite the limitations of the above points, AI is well-poised to address many current needs in medicine and the life sciences. We believe the following use cases in particular will benefit from the application of AI:

- ▶ **New, different data will provide a more holistic view of patients.** Developments in deep learning especially are increasing the kinds of data that are analyzable (text, images, sounds, time-series, etc.), and an increasing focus on health data quality and interoperability will make it easier to connect different dots of a patient's journey together into a multidimensional picture of health.
- ▶ **Algorithms will triage patients into risk groups, thereby more efficiently allocating care and augmenting provider judgment.** FDA approvals for diagnostic algorithms will increase, and these tools will optimize the medical care supply chain by routing patients to the appropriate level of treatment and specialization; furthermore, they will interface with EHRs to highlight particular risk areas, help providers follow ever-evolving standards of care and place importance on specific areas of preventive care that could prevent especially likely future illness.
- ▶ **Old drugs will be repurposed for new indications.** The explosions in granular data on patient phenotypes and our increased understanding of microscopic biochemical pathways will help life scientists revisit existing therapies in the context of new diseases to use old drugs in new ways.
- ▶ **New, more precise medications will come to market more quickly and cheaply.** The increased resolution in patient biomarkers will enable drugs to be developed for specific disease state subpopulations, and improvements in clinical trial design will push these new interventions to the clinic more easily.
- ▶ **Pharmaceutical interventions will be judged by their effectiveness, not their efficacy.** With more specific trials being undertaken and more real-world evidence being available, the success of a treatment will be determined not by whether a trial can be designed to demonstrate its effects, but rather by whether its introduction to the market positively impacts population health.





For more information, please contact:

Mathias Bellaiche, PhD

Senior Consultant, Advanced Analytics and Intelligent Automation
(479) 409-5745
mabellaiche@guidehouse.com

Acknowledgements

The author would like to acknowledge the following people for their contributions: Sri Iyer, Bassel Haidar, Alex Gromadzki, Kate Pokrass, Samuel Soltoff, and Winnie Fan.

