

# A population health perspective on artificial intelligence

Maxime Lavigne, BEng, MSc<sup>1,2</sup>; Fatima Mussa, MPH<sup>3</sup>;  
Maria I. Creatore, MSc, PhD<sup>3,4</sup>; Steven J. Hoffman, JD, PhD, LL.D.<sup>3,5,6</sup>;  
and David L. Buckeridge, MD, PhD<sup>1,2</sup>

Healthcare Management Forum  
2019, Vol. 32(4) 173-177  
© 2019 The Canadian College of  
Health Leaders. All rights reserved.



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0840470419848428

journals.sagepub.com/home/hmf



## Abstract

The burgeoning field of Artificial Intelligence (AI) has the potential to profoundly impact the public's health. Yet, to make the most of this opportunity, decision-makers must understand AI concepts. In this article, we describe approaches and fields within AI and illustrate through examples how they can contribute to informed decisions, with a focus on population health applications. We first introduce core concepts needed to understand modern uses of AI and then describe its sub-fields. Finally, we examine four sub-fields of AI most relevant to population health along with examples of available tools and frameworks. Artificial intelligence is a broad and complex field, but the tools that enable the use of AI techniques are becoming more accessible, less expensive, and easier to use than ever before. Applications of AI have the potential to assist clinicians, health system managers, policy-makers, and public health practitioners in making more precise, and potentially more effective, decisions.

## Introduction

The field of Artificial Intelligence (AI) is nearly ubiquitous with the widespread adoption of products such as automated translation, face recognition, and semantic searching. Over the last 50 years, AI-based systems have repeatedly crossed the boundary of what was thought possible with advances such as chess, automated translation, and self-driving cars. Current AI-based systems now demonstrate an unprecedented depth of reasoning and grasp of our culture.

Although AI has existed for some time, a number of factors have converged to allow the adoption of AI across a wide range of fields such as law, political science, policy, and health.<sup>1</sup> As an example of this newly sparked interest, a prominent textbook on the subject claims that AI is regularly cited as the “field I would most like to be in” by scientists in other disciplines.<sup>2</sup>

In health, the application of AI has led to improvements in many areas, such as in research on genetics<sup>3</sup> and drug discovery<sup>4</sup> and in clinical care through prevention,<sup>5</sup> diagnosis,<sup>6</sup> therapy planning,<sup>7</sup> and optimizing care delivery,<sup>8</sup> increasingly within the context of personalized medicine.<sup>9</sup> Part of the appeal of AI for health applications is the capacity of these methods to support decisions based on voluminous, heterogenous, and noisy data and to contribute to sensemaking by providing new ways to infer knowledge and relationships from data.<sup>10</sup> Despite the rapid uptake of AI in research and clinical care, the adoption has been slower in population health settings, such as for managing health systems and delivering public health.

To apply AI in population health settings in a manner that is likely to be effective, decision-makers must have some fluency in AI concepts, methods, and tools, but this can be difficult to achieve as the overarching term “AI” is used to mean many different things. To address this perceived knowledge gap, in this article, we seek to delineate different approaches and

fields within AI and illustrate through examples how different sub-fields of AI can contribute to more informed decisions, with a particular focus on current and potential population health applications. We begin by introducing the field of AI in general and presenting the core concepts needed to understand its contemporary uses. We then proceed to describe the sub-fields of AI. Finally, we examine more closely the four sub-fields of AI most relevant to population health and consider the advantages of each approach along with examples of available tools and frameworks.

## Defining the field of artificial intelligence

Any definition of AI is bound to seem recursive as it is intrinsically dependent on definitions of intelligence. Two common definitions are defining AI as “the study of the design of intelligent agents,”<sup>11</sup> and “the art of creating machines that perform functions that require intelligence when performed by people.”<sup>12</sup>

There are two key concepts embodied in these definitions. First, AI relates to intelligence that is not natural but constructed.

<sup>1</sup> Surveillance Lab, McGill Clinical and Health Informatics, McGill University, Montreal, Quebec, Canada.

<sup>2</sup> Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, Quebec, Canada.

<sup>3</sup> CIHR Institute of Population and Public Health, Canadian Institutes of Health Research, Toronto, Ontario, Canada.

<sup>4</sup> Dalla Lana School of Public Health, University of Toronto, Toronto, Ontario, Canada.

<sup>5</sup> Global Strategy Lab, York University, Toronto, Canada.

<sup>6</sup> Dahdaleh Institute for Global Health Research, Faculty of Health and Osgoode Hall Law School, York University, Toronto, Ontario, Canada.

### Corresponding author:

David L. Buckeridge, Surveillance Lab, McGill Clinical and Health Informatics, McGill University, Montreal, Quebec, Canada.

E-mail: david.buckeridge@mcgill.ca

**Table 1.** Skills required to pass the Turing test<sup>2</sup>

<i>Natural language processing (NLP) to enable communication with an interlocutor</i>	<i>Machine learning (ML) to adapt to new circumstances and detect or extrapolate patterns</i>
<i>Knowledge representation to store what is known or sensed</i>	<i>Computer vision to perceive objects and understand images</i>
<i>Automated reasoning to use stored information to answer questions and draw conclusions</i>	<i>Robotics to manipulate objects and move about</i>

Second, we cannot define AI without first agreeing on a functional definition of what intelligence is and how to recognize it.

Defining intelligence has proven to be controversial and challenging. While the debate continues, most acknowledge that humans are intelligent beings. This means that, to some extent, acting intelligently can be defined as acting in a way that is indistinguishable from how a human would act. This generally accepted criterion has been named the Turing test and passing this test was hypothesized to require the skills listed in Table 1.

In its contemporary use, advances in AI usually refer to research done in any of these six fields. The preferred term for research on artificial systems able to perform any intellectual tasks is “Artificial General Intelligence” or “Strong AI.” Another useful distinction in AI is between connectionist and symbolic approaches, which refers to the types of assumptions made and the kinds of problems addressed.<sup>13</sup>

Connectionist AI uses data-oriented approaches to derive decisions based on prior experience. The strengths of these approaches are that they are tolerant to uncertain or erroneous data and will work even if there is incomplete or no a priori knowledge about a situation. The weaknesses, however, are that the processes used are often opaque and difficult to understand and they require large amounts of data. Connectionism is akin to inductive reasoning, or a bottom-up strategy, as it learns by examples and uses prior experiences when evaluating new decisions. As an example, given a large dataset with annotated pictures of footballs and other random objects, a system can learn to classify whether the subject of a picture is a football or not. The classification is based solely on how similar or dissimilar certain features of an image are to images seen previously.

Symbolic AI uses logic and symbolic approaches to derive decisions based on an interconnected web of concepts and relations. The strength of these approaches is that they do not require data and can make decisions that are context independent. The reasoning process can also be explained and communicated. On the other hand, these approaches require the prior existence of extensive and consistent machine-readable knowledge. Symbolic approaches traditionally have also had difficulty dealing with situations that are probabilistic or continuous in nature. Symbolic AI is similar to deductive reasoning, or a top-down strategy, since it uses prior knowledge to infer and reason about new situations. Using the previous example, symbolic AI systems could be asked to infer whether a

new sport item is a football and to explain what is missing or what information is not relevant.

It has been proposed that both connectionist and symbolic AI approaches are needed to achieve human-like cognition.<sup>13</sup> The degree to which each approach is used within an application depends on the task at hand. Different sub-fields of AI tend to favour one of the two approaches. For example, knowledge representation tends to be mostly symbolic, while machine learning is mostly connectionist, and automated reasoning tends to use both approaches. Below we consider in greater detail the four sub-fields of AI particularly relevant to population health: Natural Language Processing (NLP), knowledge representation, automated reasoning, and machine learning. In the following sections, we first present a short summary of each field, followed by a description of recent advances relevant to population health, and finally some examples of real-world applications, which interested readers can use to explore the field further.

## Natural language processing

The field of NLP studies human-computer interactions made through natural human languages, such as English. In other words, it aims to allow computers to (1) interpret natural languages, (2) generate responses using natural languages, and (3) learn new concepts and relations from interactions using natural language.<sup>2</sup> Advances in NLP have led to the development of frameworks and resources that are now easy to use, readily available, and often free of charge. One such example is spaCy, which is described as “industrial-strength natural language processing” using the Python programming language.<sup>14</sup>

This field also includes voice recognition, with some commercial examples of this technology being stable enough that they are currently sold as digital assistants. Major voice recognition platforms often offer programmatic interfaces allowing their use by third parties and researchers.

In health applications, NLP tools are useful for extracting data from patient medical files, published articles, or even from social media. Once text is extracted, NLP methods can also be used to identify the context and meaning of communication. Examples include case detection<sup>15</sup> or phenotyping<sup>16</sup> from electronic medical record data, quality assurance,<sup>17</sup> and analyzing social media data to assess vaccine hesitancy,<sup>18</sup> and detect and track emerging health threats.<sup>19</sup>

## Knowledge representation and automated reasoning

We consider these sub-fields together because automated reasoning often makes extensive use of knowledge representation. “Automated reasoning” occurs when stored knowledge is used to answer questions and to draw new conclusions, while “knowledge representation” refers to the encoding of human knowledge in a way that is machine-readable, clear, unambiguous, and consistent.<sup>2</sup>

Knowledge representation often makes use of formal, explicit specifications of a shared conceptualization, called

software ontologies,<sup>20</sup> that describe relations, properties, and categories of concepts and entities. Ontologies usually have well-defined domains and can be used for applications such as creating collaboratively editable representations of domain knowledge. The languages currently used to encode knowledge formally in ontologies, however, are limited in their ability to represent some types of knowledge, such as temporal and complex relationships, which makes certain problems intractable.<sup>21</sup> The ICD-11 project and “IBM Watson” are two well-known examples of this field of AI, with the first encoding knowledge about diseases and the second using reasoning and deduction in order to win trivia games and more recently to guide clinical diagnostic and therapeutic decisions. Another example of greater relevance to population health is the use of knowledge representation to encode practice guidelines.<sup>22</sup> This application of AI ensures that guidelines can be easily updated and verified, and any improvements can be automatically applied to every system that incorporates this knowledge. Such an approach has the potential to save time and effort in addition to creating a safer environment in which decisions are made based on consistent, coordinated, and up-to-date information that reflects input and updates from all relevant collaborators.

Automated reasoning allows researchers to study decision-making under constraints and provides a foundation for decision support systems. The aim of automated reasoning is to develop methods that can propose logically or probabilistically sound solutions given a specific question and context. Automated reasoning is therefore well-positioned to support evidence-based decision-making. It can, for example, support decision-making and planning for interventions, which are integral to both clinical and population health.

An example of the use of automated reasoning could be a system that chooses the most appropriate intervention given what is known about a specific patient or population. Implementation of this type of decision support is critical for realizing the promise of precision medicine in a clinical context.<sup>23</sup> In a population health context, automated reasoning plays a similarly central role in supporting the implementation of learning public health systems, particularly in knowledge translation<sup>24</sup> and precision public health.<sup>25</sup>

## Machine learning

Machine learning has been defined as “the study of data-driven methods capable of mimicking, understanding and aiding human and biological information processing tasks. [...] In the broadest sense, machine learning and related fields aim to ‘learn something useful’ about the environment within which the agent operates”.<sup>26</sup> There are many similarities between statistics, commonly used in health applications, and machine learning. At a general level, both are approaches to learning from data, although traditional statistics assumes observed data are generated by a probabilistic data model, while machine learning assumes data are generated through an unknown mechanism. This difference between the approaches has led to some well-known criticisms of traditional statistics; for

example, some have argued that “by being committed to the almost exclusive use of data models, this commitment has led to irrelevant theory, questionable conclusions, and has kept statisticians from working on a large range of interesting current problems.”<sup>27</sup>

One of the major differences between the two paradigms is the underlying objective of the analysis. In machine learning, the main aim is usually predictive accuracy or generating predictions that closely match observed data. In contrast, in traditional statistics, the aim is usually to find the model that best represents the data. Once a model is defined and parameters are estimated, any conclusions made are conditional on how appropriately the model represents the true underlying process. As increasingly more complex data models are proposed and used, proponents of machine learning suggest that automated approaches to model-building should be considered.<sup>27</sup> However, by moving away from parametric models, there is a chance to lose sight of the original purpose of the inquiry. This situation can be particularly problematic in fields such as epidemiology, where questions often relate to the data generation mechanism itself (ie, identifying causal mechanisms) and are less frequently about how closely an action can be replicated (ie, predicting likely outcomes). However, some machine learning researchers have suggested that traditional statistical methods may also not be that useful for answering causal questions. The reason being that the natural mechanisms of interest are often either unobserved or unobservable and therefore estimates of their properties may be artificial products of the assumptions made in specifying the data model.<sup>27</sup>

An important classification of machine learning approaches is based on the type of learning used. The three major types of learning are supervised, unsupervised, and reinforcement learning.<sup>2</sup> In supervised learning, agents are provided with input-output pairs that have been correctly identified beforehand. Here, agents learn to replicate or mimic the correct answer as closely as they can. Supervised learning is used in tasks such as prediction and classification. In unsupervised learning, however, agents learn patterns and identify relations from data without explicit labelling of the outcome. This approach is useful for applications such as clustering, pattern recognition, or the discovery of latent factors. Finally, in reinforcement learning, agents learn from positive or negative interactions with their environment and themselves. The kinds of tasks being achieved will depend on the reinforcement provided.

Both statistical learning and machine learning require large datasets of good quality. Any model that learns from data will be improved by more up-to-date, diverse, expressive, and larger amounts of data. Similarly, poor quality data, such as data sampled in a biased manner, can result in models that make erroneous or biased predictions. Consequently, methodologies aimed at improving the process of gathering and analyzing data will have a strong impact on the performance of machine learning algorithms. This relationship between machine learning and “Big Data” is why the two topics are often discussed in the same context. Many of the recent

advances in AI applied to health are from the use of machine learning, and in particular the use of deep neural networks for supervised learning. For example, these types of AI-based approaches have been shown to classify the results of radiological imaging more accurately than human experts.<sup>28</sup> This level of performance has yet to be demonstrated in supporting high-level decision-making in population health settings, possibly because the data tend to be more complex in structure. For example, because concepts are often measured indirectly, it can be challenging to ensure the data represent the concept in an unbiased manner, and correct interpretation requires a broad set of domain and contextual knowledge. Some promising applications have been reported, however, in areas such as predicting population characteristics from remote sensing data<sup>29</sup> and predicting environmental exposures from historical satellite/street-level images.<sup>30</sup>

## Conclusion

In this article, we described and attempted to delineate different approaches and fields within AI to help cut through some of the hype and jargon and to assist readers in identifying opportunities for AI in population health. We also considered in greater detail the four sub-fields of AI that we believe have the greatest potential for population health applications.

A central message of this article is that AI is a broad field with many sub-fields, even though AI is often used to refer to one sub-field or approach, such as knowledge representation or machine learning. The sub-field of machine learning in particular is now being applied to many clinical problems, but machine learning has been applied less frequently to population health or public health problems. Machine learning approaches are similar and complementary to traditional statistical approaches, and both have the potential to play an important role in using “Big Data” to understand and predict healthcare and population health outcomes. However, particularly when applying these approaches to decision-making or predictions at a population level, attention must be paid to the potential for these approaches to produce health inequities, either through the use of biased data or through uneven access to the technology. Predictions and models based on non-representative or biased data can propagate underlying biases and exacerbate health inequities at a population level if sufficient care is not taken to mitigate these issues.

As a practical example to illustrate the current potential of AI, it would be possible to develop an automated system that uses NLP to read medical charts to detect cases of a disease of public health importance from plain text. These new cases could be linked to current knowledge using the ICD-11 disease ontology. Once resolved to known conditions, the cases could be analysed by machine learning methods in a surveillance system to detect outliers or clusters. Unusual patterns could then be presented to practitioners, drawing on other knowledge to highlight the potential magnitude of a public health problem and to predict the potential growth in cases. Finally, an expert system could then integrate the information about the cases


with computable knowledge about potential interventions and use automated reasoning to propose effective public health interventions.


While predicting the future is always challenging, it appears that the contributions of AI to clinical and population health will continue to expand. The tools that enable the use of AI techniques are becoming more accessible, less expensive, and easier to use than ever before. Applications of AI have the potential to assist clinicians, health system managers, public health practitioners, and policy-makers in making more precise, and potentially more effective, decisions. Efforts to do so, however, are more likely to be successful if they are developed through collaborations across sectors and if they make available interoperable AI-based applications that use domain knowledge and learn from experience.

## Declaration of conflicting interests

M.L. and D.L.B. work for McGill University. F.M., M.I.C., and S.J.H. work for the Canadian Institutes of Health Research (CIHR), the Government of Canada’s national health research investment agency. The views expressed in this article are those of the authors and do not necessarily reflect those of CIHR or the Government of Canada.

## ORCID iD

Maxime Lavigne  <https://orcid.org/0000-0002-6401-6381>

David L. Buckeridge  <https://orcid.org/0000-0003-1817-5047>

## References

1. Stone P, Brooks R, Brynjolfsson E, et al. Artificial Intelligence and Life in 2030. *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*. Stanford University, Stanford, CA; September 2016. <http://ai100.stanford.edu/2016-report>
2. Russell SJ, Norvig P. *Artificial Intelligence: A Modern Approach*. Harlow, Malaysia: Pearson Education Limited; 2016.
3. Libbrecht MW, Noble WS. Machine learning applications in genetics and genomics. *Nat Rev Genet*. 2015;16(6):321-332.
4. Fleming N. How artificial intelligence is changing drug discovery. *Nature*. 2018;557(7707):S55-S57.
5. Liu Y, Kohlberger T, Norouzi M, et al. Artificial intelligence—based breast cancer nodal metastasis detection: insights into the black box for pathologists. *Arch Pathol Lab Med*. 2018.
6. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316(22):2402-2410.
7. Babier A, Boutillier JJ, McNiven AL, Chan TCY. Knowledge-based automated planning for oropharyngeal cancer. *Med Phys*. 2018;45(7):2875-2883.
8. Bates DW, Saria S, Ohno-Machado L, Shah A, Escobar G. Big data in health care: using analytics to identify and manage high-risk and high-cost patients. *Health Aff*. 2014;33(7):1123-1131.
9. Fröhlich H, Balling R, Beerenwinkel N, et al. From hype to reality: data science enabling personalized medicine. *BMC Med*. 2018;16(1):150.

10. Bellazzi R, Zupan B. Predictive data mining in clinical medicine: current issues and guidelines. *Int J Med Inform.* 2008;77(2): 81-97.
11. Poole DL, Mackworth AK, Goebel R. *Computational Intelligence: A Logical Approach*. New York, NY: Oxford University Press New York; 1998. Vol. 1.
12. Kurzweil R. *The Age of Intelligent Machines*. Cambridge, MA: MIT Press; 1990.
13. Minsky ML. Logical versus analogical or symbolic versus connectionist or neat versus scruffy. *AI Magazine.* 1991;12(2):34.
14. Honnibal M, Montani I. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. 2017.
15. Brownstein JS, Freifeld CC, Madoff LC. Digital disease detection—harnessing the web for public health surveillance. *N Eng J Med.* 2009;360(21):2153-2157.
16. Richesson RL, Sun J, Pathak J, Kho AN, Denny JC. Clinical phenotyping in selected national networks: demonstrating the need for high-throughput, portable, and computational methods. *Artif Intell Med.* 2016;71:57-61.
17. Freifeld CC, Brownstein JS, Menone CM, et al. Digital drug safety surveillance: monitoring pharmaceutical products in twitter. *Drug Saf.* 2014;37(5):343-350.
18. Salathé M, Khandelwal S. Assessing vaccination sentiments with on-line social media: implications for infectious disease dynamics and control. *PLoS Comput Biol.* 2011;7(10):e1002199.
19. Mawudeku A, Blench M. Global public health intelligence network (GPHIN). Paper presented at: 7<sup>th</sup> Conference of the Association for Machine Translation in the Americas, Cambridge, Massachusetts; August 8-12, 2006;8-12.
20. Staab S, Studer R. *Handbook on Ontologies*. Berlin, Heidelberg: Springer Science & Business Media; 2010.
21. Lukasiewicz T, Straccia U. Managing uncertainty and vagueness in description logics for the semantic web. *Web Semant.* 2008; 6(4):291-308.
22. Peleg M. Computer-interpretable clinical guidelines: a methodological review. *J Biomed Inform.* 2013;46(4):744-763.
23. Sittig DF, Wright A, Osheroff JA, et al. Grand challenges in clinical decision support. *J Biomed Inform.* 2008;41(2):387-392.
24. Riaño D, Real F, López-Vallverdú JA, et al. An ontology-based personalization of healthcare knowledge to support clinical decisions for chronically ill patients. *J Biomed Inform.* 2012; 45(3):429-446.
25. Haendel MA, Chute CG, Robinson PN. Classification, ontology, and precision medicine. *New Eng J Med.* 2018;379:1452-1462.
26. Barber D. *Bayesian Reasoning and Machine Learning*. New York, NY: Cambridge University Press; 2012.
27. Breiman L. Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Stat Sci.* 2001;16(3):199-231.
28. Yu KH, Zhang C, Berry GJ, et al. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. *Nat Commun.* 2016;7:12474.
29. Gebru T, Krause J, Wang Y, et al. Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States. *Proc Nat Acad Sci.* 2017;114(50):13108-13113.
30. Weichenthal S, Hatzopoulou M, Brauer M. A picture tells a thousand... exposures: opportunities and challenges of deep learning image analyses in exposure science and environmental epidemiology. *Environ Int.* 2019;122:3-10.