

Tutorial

An Introduction to AI for Clinicians: Tutorial

Stephen B Lee¹, MS, MD; Alexis B Carter², MD; Muhammad Hamis Haider³, PhD; Seok-Bum Ko³, PhD

¹Division of Infectious Diseases, University of Saskatchewan, Regina, SK, Canada

²Department of Pathology and Laboratory Medicine, Emory University, Atlanta, GA, United States

³Department of Electrical and Computer Engineering, University of Saskatchewan, Saskatoon, SK, Canada

Corresponding Author:

Stephen B Lee, MS, MD

Division of Infectious Diseases

University of Saskatchewan

1440-14th Avenue, Regina General Hospital, 2nd Floor Medical Office Wing, ID Clinic

Regina, SK S4P 0W5

Canada

Phone: 1 3067664247

Email: leestephenz@gmail.com

Abstract

Artificial intelligence (AI) is already fundamentally changing society, with medicine being no exception. AI will impact how we practice, how hospitals operate, and even the practice of medicine itself. The use of AI-based products has already begun, with examples including AI scribes and large language models such as ChatGPT. Work is ongoing to produce models that have specific functions within medicine, such as kidney injury prediction. However, transformative foundational work, such as AlphaFold (for protein structure prediction), also promises to completely change the way we approach medicine. Therefore, clinicians must develop a clear understanding of AI, not as an optional skill, but as a core competency of modern medical practice. This paper serves as a tutorial to guide medical professionals through the basic principles of AI. It will teach clinicians how to build a mental scaffold to understand and springboard into AI. The core parts of this paper are organized in steps, with additional relevant topics addressed in modules at the end of the paper. The core steps are meant to be read sequentially. To prepare the reader for the rest of the paper, this tutorial will first introduce what AI is and then cover some basic definitions needed to understand other concepts. The reader will then be ready to understand what deep learning is and the difference between supervised and unsupervised learning. Finally, the reader will go through how deep learning models learn. Separate modules on safety and clinical applications are also included. This tutorial is relevant to clinicians at all levels but may be particularly useful for practicing clinicians who are encountering AI tools integrated into their practices without previous formal education in the field. Users of this tutorial can refer to specific sections or read the entire paper.

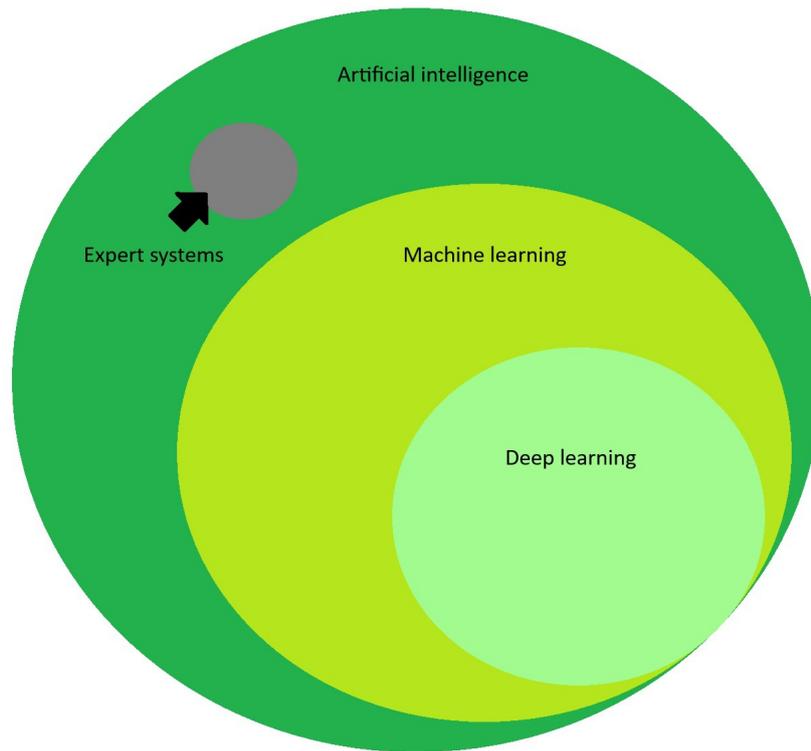
Interact J Med Res 2026;15:e85266; doi: [10.2196/85266](https://doi.org/10.2196/85266)

Keywords: medical education; machine learning; artificial intelligence; AI; deep learning

Step 1: Understanding What Artificial Intelligence Is

Before delving into the concepts underlying artificial intelligence (AI), it is important to understand what AI means. AI is a broad term that generally refers to computer

systems that can perform complex tasks historically associated with humans, such as human learning, comprehension, problem-solving, decision-making, creativity, and autonomy [1]. In computer science, AI algorithms encompass a variety of methodologies. Most AI in the modern era is machine learning (ML) and even more specifically, deep learning (DL; [Figure 1](#)). While all ML is considered AI, not all AI is ML.

Figure 1. Venn diagram of artificial intelligence terminology.

Expert systems are an older paradigm of AI in which a subject matter expert's knowledge is hardcoded into intricate rule-based algorithms to simulate human decision-making (eg, "if x condition, then y result"). Well-known examples include the MYCIN system, designed by Edward Shortliffe in 1974 to predict antibiotic choice [2], and traditional chess AI systems. Even famous examples such as Deep Blue (IBM Corp) relied on hardcoded knowledge [3]. Expert systems and these hardcoded methodologies have been largely abandoned in the modern era due to the incredible effort required to develop them and the brittleness and inflexibility of the resultant systems. Limited examples of expert systems also exist in health care, such as older sepsis and drug interaction clinical decision support systems. AI is primarily ML currently. Rather than relying on hardcoding from experts, ML algorithms are trained to learn relationships and patterns in data. The terms DL and ML are often used interchangeably; however, it is worthwhile mentioning that DL is really a subset of ML. There are many types of ML that are not DL (eg, k-means clustering and decision trees). DL is discussed in depth in a following section.

Step 2: Understanding Basic Definitions

Some level-setting is required to understand ML algorithms and models. ML *algorithms* are tools (eg, logistic regression) that are trained on data to create an ML *model*. ML models that are used for classification are frequently referred to as *classifiers*. The quality of the data used to train an algorithm is critical to the performance of the model. Each record or event in the dataset is referred to as an *instance*. Each aspect of an instance (eg, color, duration, and test result) used to

train a model is known as a *feature*. In a simple set of data, instances would be rows in your data spreadsheet, while features would be the column headers. Many ML datasets can have thousands of instances and hundreds of features.

A *label* is the information on a particular feature for a certain instance in the dataset (eg, "red" for a feature of "color," "30 minutes" for a feature of "duration," and "34 mg/dL" for a feature of "glucose level"). A label generally refers to information that has been applied by a human or another algorithm based on manual or traditional analysis of the data and represents a classification, categorization, ranking, or answer to a question. For example, in a dataset of inpatients hospitalized for at least 30 days with a feature of "patient developed a hospital-acquired infection during admission," the label would be "positive" if an infection was acquired and "negative" if none occurred, bearing in mind that the application of positive or negative would have to be done by a human who was analyzing the patient's data.

Of note is that data that are not manually labeled are simply referred to as data. One may also choose to use only a subset of the data in a dataset for training.

Step 3: Differentiating Between Supervised and Unsupervised ML

A core initial concept is understanding the difference between 2 major categories of ML. In general, ML can be classified into supervised and unsupervised learning, although there are other categories or approaches such as reinforcement and transfer learning.

In *supervised learning*, data that have been labeled are fed into the ML algorithm for training. Once the algorithm has been trained, it is called a *model*. One of the greatest challenges in developing a high-quality model is being able to obtain a substantial amount of data (ie, usually thousands of instances) that have been accurately labeled. An example of supervised learning would be an algorithm trained

on a dataset of chest X-rays and corresponding diagnoses [4-6]. In *unsupervised learning*, unlabeled data are fed into an algorithm, and the algorithm discovers relationships and groupings for itself. This difference is illustrated in Figure 2. An example of unsupervised learning is work in which an algorithm identified distinct clusters or groups of patients who had COVID-19 [7,8].

Figure 2. Illustration of (A) supervised and (B) unsupervised learning.

Index	Data	Label
Data point 1		Circle
Data point 2		Square
Data point 3		Circle
Data point 4		Square
Data point 5		Square
Data point 6		Circle

Figure 2(A)

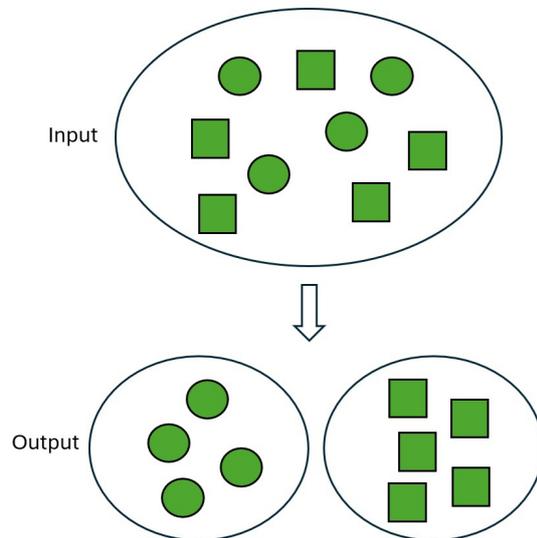


Figure 2(B)

In Figure 2A, the shapes are labeled as circles and squares by a human. The dataset is fed into the algorithm, and through these labels, the machine learns which features contribute most to classifying the shapes. In Figure 2B, the machine is fed raw data without any labels. Through exploration and differences detected in the data among its features, the model learns that there are potentially 2 different categories of objects. Of note is that the model may not inherently recognize them as a circle or a square but rather as 2 distinct categories of objects.

Other forms of learning also exist, such as *reinforcement learning*. In reinforcement learning, an algorithm experiences an environment and takes an action. On the basis of this action, the model is given a reward or a punishment as feedback. The algorithm learns which patterns result in rewards vs punishments and adjusts its behavior accordingly. Furthermore, supervised and unsupervised learning exist on a continuum, and some forms of learning are a mixture of both (ie, some instances are labeled, while others are not). A full description of these concepts is beyond the scope of this paper.

Transfer learning is another method in which an algorithm is first trained on a very large but nonspecific set of data for the desired outcome and then uses the learned patterns on a smaller but more specific set of data to refine the algorithm into a model. Transfer learning is often used where high-quality labeled datasets specific to the subject area are limited, whereas broader, nonspecific datasets are more plentiful.

Step 4: Defining DL and Neural Networks

A subcategory of ML is DL. This branch of ML is heavily focused on neural networks. *Neural networks* were first described in the mid-20th century [9,10] and were designed to emulate neural processes in the human nervous system. While foundational work has been ongoing for decades in the field [11-13], much early work was constrained by hardware and data availability. The emergence of powerful parallel computing, called graphics processing units; platforms to leverage graphics processing units; and the availability of large datasets have helped overcome these barriers [14].

Artificial nodes, called *neurons*, are connected in layers to form a network. Data are put into the input layer of the network; the network processes the data through one to many hidden layers and then provides results in the output layer.

DL specifically refers to ML done on neural networks with many layers (thus the term “deep”). There is no exact number of layers that is generally agreed upon as a threshold. However, common examples such as ResNet (Microsoft) and the architecture underlying ChatGPT (OpenAI) can contain a hundred layers and billions of parameters [15,16].

Step 5a: Walk-Through on How ML Algorithms Work

Next, this tutorial will lay out the process by which ML algorithms work, describe the process with an example, and then define other key terms. To learn how to properly predict, classify, or rank instances in supervised learning, an algorithm analyzes the data to determine which features contribute the most to the data labels. This learning process is called *training*. During training, the machine learns to adjust internal parameters, called *weights*, to produce the desired outputs. These weights correspond to individual neurons. This occurs through minimizing loss functions, in which the gap between the observed and expected prediction is minimized.

For example, consider a model designed to estimate the risk of candidiasis in hospitalized patients. The model may learn that features such as intensive care admission, fever, abdominal symptoms, or a normal white blood cell count are associated with either a higher or lower risk of candidiasis. Features will often be mapped in a complex fashion across a series of neurons and layers. Modification of the weights applied to each of these features influences the resulting model's output through complex interactions across many layers of analysis, allowing the model to learn patterns that relate input features to the target outcome. In simpler forms of ML, it is sometimes possible to determine which features most greatly contribute to an outcome; however, in DL, this is often difficult because of the number of nodes and layers involved. Therefore, in DL, the change in importance may not be directly interpretable.

Datasets are often large and contain many features, some of which logically have nothing to do with the outcome being examined. If the data used to train an algorithm contain these completely unrelated variables, the algorithm may make nonsensical associations, which can later result in spuriously wrong results. For example, it may associate the color of a hospital gown and/or patient's sandwich preference with overall length of stay. However, if there are different colors of hospital gowns or menu choices for patients in intensive care units vs regular floors that are not captured in the data, then the model may miss the confounding variable completely. These errors in model development are hard to detect in all ML and even harder in DL.

Step 5b: Walk-Through of Basic ML Mathematics Using an Example – Loss Function

To illustrate the mathematics of ML training, we can use an extremely simplified example with a continuous variable in supervised learning. It is noted that most modern DL requires massive sets of data. While specifics may change in other forms of ML, this example will illustrate general concepts. In this example, the dataset includes the trough levels of a nephrotoxic drug and the resultant measured (true) estimated glomerular filtration rate (eGFR) in patients. The measured eGFR is the labeled data in this supervised ML. An appropriate ML algorithm is selected that attempts to predict the eGFR based on drug levels (Table 1).

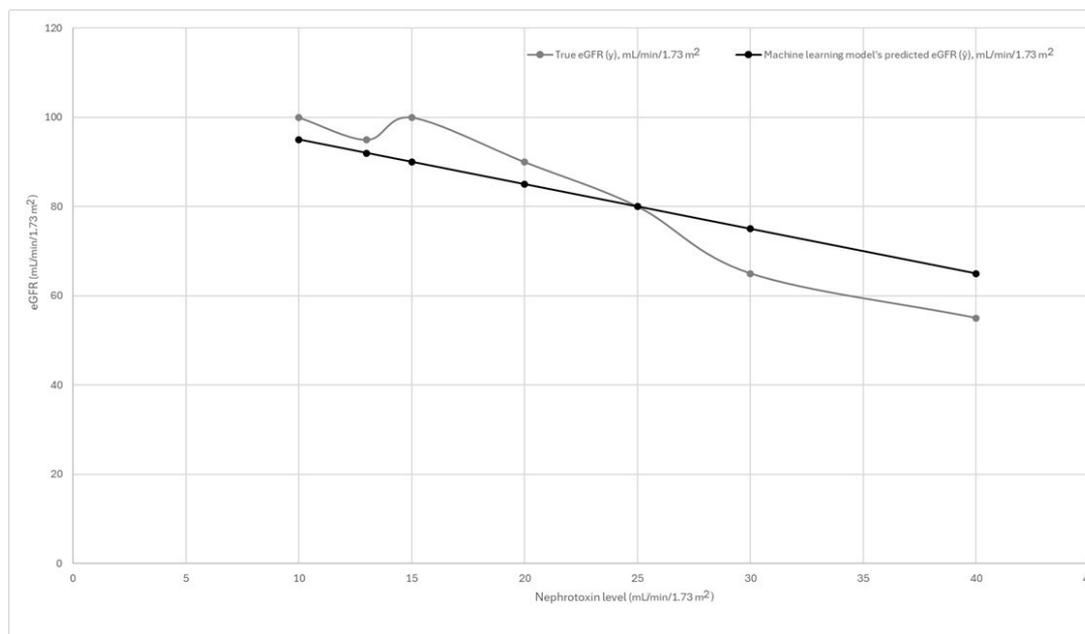
Table 1. Sample data for a model predicting nephrotoxin toxicity.

Nephrotoxin level (mg/L)	True eGFR ^a (y; mL/min/1.73 m ²)	Machine learning model's predicted eGFR (\hat{y} ; mL/min/1.73 m ²)
10	100	95
13	95	92
15	100	90
20	90	85
25	80	80
30	65	75
40	55	65

^aeGFR: estimated glomerular filtration rate.

The model produces a prediction (\hat{y}) based on the nephrotoxin level (x). The algorithm compares the predicted eGFR (\hat{y}) to the real measured value of eGFR (y). Graphically, the predicted values of eGFR (the orange line in Figure 3) are plotted against the true values (the blue line). To determine the pattern between the nephrotoxin level (x) and

the measured value of eGFR (y), the algorithm calculates the loss function. Differences might appear small in most regions but noticeable in some. To capture a model's performance across all data points, we use this loss function, a mathematical function that quantifies the total error between predicted and observed values.

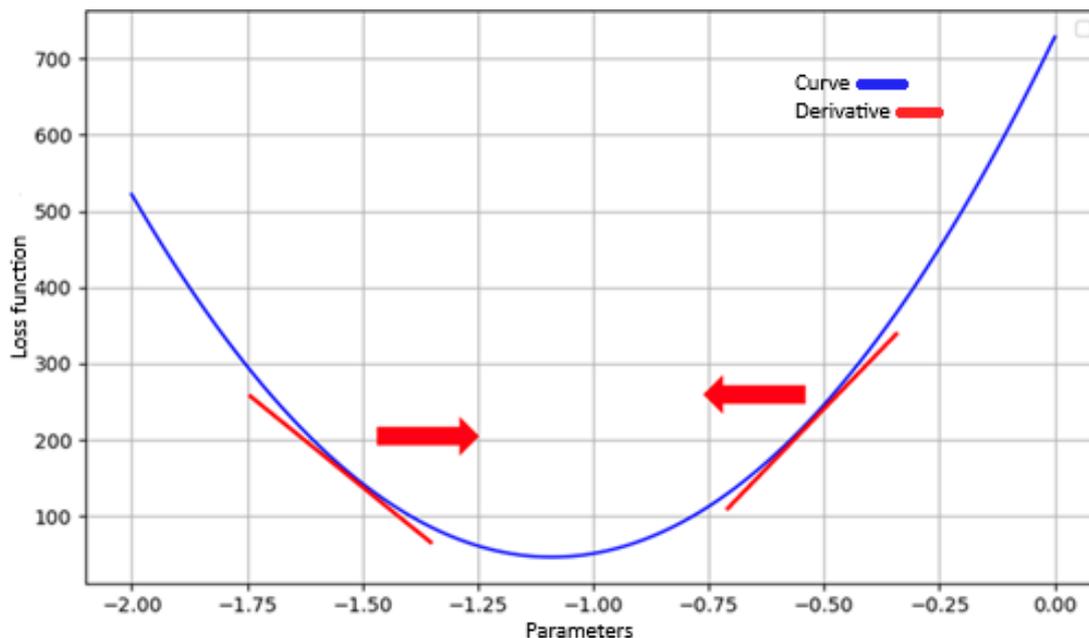
Figure 3. Sample data for a model predicting nephrotoxin level. eGFR: estimated glomerular filtration rate.

By convention, updates are made to the loss function after each batch in training. The model then adjusts the weights and tries again to determine, across all the instances in the dataset, which set of weights will produce the lowest loss function on average for the training data. With high-quality data that represent the full spectrum of possible cases that the model could be given, this helps ensure that the model is trained to perform well on average and does not result in spuriously wrong predictions. The specifics of how the algorithm calculates loss functions vary between algorithm types, with different functions being optimized for different tasks. For instance, the mean squared error, a common loss function, squares the difference between \hat{y} and y . By squaring the difference, it ensures negative loss values do not cancel positive values when summed, and it helps penalize severely wrong predictions.

In a model that performs well when it is fed new data (ie, generalizes well), \hat{y} and y will be similar for any new patient's nephrotoxin level fed into it. Models that do not generalize well may make large, nonsensical errors in prediction for a small subset of patients with a few specific differences in data points. This is why training a model on high-quality data that accurately represent the types of data it may encounter after deployment is critical.

Step 6: Understanding the Concepts of Backpropagation and Gradient Descent

Another set of key concepts to understand is backpropagation and gradient descent. In DL, a model modifies the weights of specific neurons within a neural network to create its predictions. When data are input, they forward pass through the network. The initial run creates random weights, and the prediction ability of the initial model is likely poor. However, after this forward pass, the model evaluates its resultant loss function and attempts to minimize the loss through a process called gradient descent. In gradient descent, we assign a learning rate, which increments our point along the curve of the loss function (Figure 4). If the model discovers that the loss is increasing, it will move backward to decrease the loss. Conversely, if it discovers that the loss is decreasing, it will continue moving in that direction. In doing so, it eventually seeks out the local minima, thus minimizing the loss function and improving the model's predictive capacity.

Figure 4. Graphical illustrations of concepts.

Mathematically, gradient descent can be expressed as $w = w - \eta \left(\frac{dL}{dw} \right)$. Remember that the derivative $\left(\frac{dL}{dw} \right)$ of the loss function gives the gradient at any given point. This allows the model to know if the loss is increasing or decreasing. The learning rate (η) tells the model how much to move in each direction. Recall that the loss function quantifies the difference between the predicted and observed values, with the goal of identifying the lowest point and, thus, minimizing this difference. These concepts are illustrated again visually in Figure 4. The mean squared error of our example model's simple line (Figure 3) is a parabola, creating an easy visualization of gradient descent for illustration purposes. When η is set at too large a value, it is possible that the model will jump very far back, such that it misses the local minima. Conversely, too small a value of η may result in such minimal movement that the minima is never reached.

During the training process, to specifically update the gradient, the model will undergo a process called backpropagation. The trainer algorithm will go through the neural network and alter the weights of neurons with the intention of decreasing loss. This process ultimately results in better prediction ability of the ML model.

Module 1: A Brief Explanation of AI Explainability and Safety

As AI becomes significantly more powerful and integrated into society, various risks and errors have been observed. A brief discussion of safety is outlined here, and there are many other publications available that dive into each of these in detail.

How an ML model determines its output (ie, resultant prediction or result given to a user) from input is often unclear. This refers to the “black-box” nature of all ML but particularly of DL and neural networks. A field called

explainable AI has emerged and attempts to not only better understand how models make their predictions but also add components to the ML model that require the model to explain how it arrived at its result (ie, which features had the most impact) [17].

For all AI, especially in medicine, it is imperative that model performance be checked for variations in outcomes that may indicate that human bias in the data has been promulgated or exacerbated by the model. AI algorithms are exquisitely sensitive pattern detection tools. As such, they can detect slight variations in data that resulted from human prejudice and bias. Worse, they can promulgate such bias into the model. Therefore, it is imperative that models be checked for differences in results that can be attributed to race, gender, socioeconomic status, religion, etc, as the presence of these elements in a model will cause the model to produce inaccurate results in certain groups of patients. Concrete examples include models having difficulty diagnosing dermatological conditions in those with darker skin tones [18], and a software program that accidentally referred White patients over African American patients to receive special care [19]. In the latter example, while the goal was to ensure patients received the care required, the algorithm was designed to predict whose care would cost more money, and it was found that less money was spent on African American patients despite having the same level of need [19]. Furthermore, a study found that an AI system could learn to predict ethnicity from radiographs alone [6]. There are now tools that assist in the detection of these elements, and models that include these elements need to be retrained on optimized data or otherwise mitigated to ensure the best care for all patients. Finally, risks associated with the infrastructure of AI also exist. As many models use a cloud-based model for computation and some public commercial large language models (LLMs) use input data to retrain models, clinicians need to be aware of where sensitive data are being sent and stored.

Now that the use of AI is rapidly becoming more pervasive, it is reasonable to think that the degree to which AI is set to perform autonomously (ie, without a human in the middle) will increase. Some publications have discussed safety measures that are required to mitigate risks as AI becomes more autonomous and integrated into systems. These risks range from tangible risks today to theoretical risks with more powerful models. Experts have classified these risks into 4 categories: misuse, misalignment, mistakes, and structural risks. Misuse occurs when users intentionally instruct an AI tool to behave in harmful ways (ie, the user is an adversary). Misalignment occurs when AI systems knowingly act against human intent (ie, the AI is an adversary). This includes intentional deception by AI. Mistakes occur when AI systems produce incorrect outputs without intentional wrongdoing, often because real-world data are complex and influenced by many contributing factors. Finally, structural risks may emerge whereby pervasive AI systems integrated into society cause harm through the actions of multiple independent agents in a multifactorial, multiagent fashion [20].

A substantial body of ongoing research in AI focuses on ensuring and improving AI safety. For instance, studies focus on the effects of adversarial attacks on models to understand safety. Developers also attempt to build safeguards into models and use red teaming, where security professionals attempt to simulate attacks to determine robustness [4,18]. Equally important are efforts to create nuanced and well-informed regulations and guidelines for development [18,21-26].

Module 2: Contemporary Clinical Applications of AI in Medicine

Health care is one of the most promising industries for AI. While disruptive-level work is underway, such as the ability to understand protein folding [27], AI has numerous applications currently being used routinely in health care.

Documentation is recognized to often be excessive and contributory to physician burnout, with an American Medical Informatics Association survey finding that 73.26% of health care professionals believed the time spent was inappropriate, 77.42% reported after-hours work related to documentation, and 74.83% believed that documentation impedes patient care [28]. Canadian data show similar findings, with physicians spending excessive amounts of time on administrative tasks that result in burnout [29,30].

AI scribes are tools that can ambiently listen to patient interactions and automatically generate notes for physicians, reducing the administrative burden for physicians. Numerous companies have created offerings; however, in general, scribes use LLMs, models based on the transformer architecture (which uses an attention mechanism to process preceding information and learn relationships among them) [31,32]. Within the context of scribes, the LLM uses this mechanism of self-attention to help generate logical text. Thus, the same limitations of transformers and LLMs carry forward onto

AI scribes. For example, in many LLMs, hallucinations are a concern, which are theorized to arise because of algorithms being rewarded for correct responses, thereby making guessing a more advantageous response than acknowledging uncertainty. Health care providers must be aware of these limitations of AI scribes and how they may potentially arise.

LLMs have also been used by vendors such as Epic to automatically extract information out of existing notes, such as creating discharge summaries, reading radiology reports, providing summaries, preparing tasks based on the note being created, and providing insights [33]. They have also been used as chatbots to act in clerical roles and as search engines for medical knowledge [34,35]. While LLMs such as ChatGPT (OpenAI), Claude (Anthropic), and Gemini (Google) are commonly used by the general public, OpenEvidence (OpenEvidence LLC) is trained specifically on medical literature, reducing errors and hallucinations [36]. Models intended for use in research and science also exist, such as Perplexity (Perplexity AI) and Elicit (Elicit Research) [36,37]. Open source, medically-focused models without interfaces also exist, such as MedGemma (Google) [38].

Under the broader definition discussed in this paper, AI has been used in diagnostics for decades using simple rule-based algorithms. However, more recently, ML-DL-based approaches have begun to gain traction due to improved performance. Areas that have shown promise for ML-DL performance are radiology and pathology [39,40]. In radiology, a DL software such as CINA-iPE (Avicenna.AI) has shown promise in detecting pulmonary embolisms [41]. In pathology, ML has improved rapid patient diagnostics based on DNA methylation markers [42]. Sepsis and cardiac arrest prediction have been an ongoing area of work [43,44]. While results from this work are promising, due in part to interpretability, accuracy, and the potential for uncovering clinically irrelevant abnormalities (“incidentalomas”), it is unclear how clinicians should react to findings [45]. The volume of AI tool uptake is increasing, and AI will inevitably impact the workflow of clinicians in the near future. AI could also support hospital infrastructure [46].

Despite its potential benefit, AI integration into health care is still an evolving landscape. Regulation and guidance remain an important area of evolving work for health care AI, with numerous national and international bodies producing frameworks and guidelines for AI use [47-49]. Important questions that remain under debate include the responsibility for AI errors, with many institutions holding physicians ultimately accountable for clinical decisions, how to mitigate bias in health care AI propagated by inherent bias in training data, and the importance of privacy [6,18,19,50]. Furthermore, many studies have indicated that implementation into workflows can also be challenging, due to a lack of awareness and engagement by both patients and health care professionals, as well as logistical implementation challenges inherent to any health technology [51].

Future Directions

AI promises to be one of the most critical revolutions of human society, arguably on par with the industrial or even agricultural revolutions. AI will impact every area of human society, including health care.

While current implementations such as LLMs, predictive models such as convolutional neural networks, and the tools they have created (eg, AI scribes and ChatGPT) are influential, most of the society-changing work is being dedicated to creating artificial general intelligence and artificial superintelligence. While the exact definition of these terms, and even their possibility, is a matter of contention [52], they generally refer to AI systems that are either as intelligent as or more intelligent than human beings across a wide array of domains and tasks.

In achieving this goal, AI will have implications on the role of humans in a post-artificial general intelligence society. While there are critical concerns about human displacement, there is also a potential for creating abundance, reducing scarcity, and an ability to supercharge scientific discovery [53].

Due to its importance, the authors believe it is important that clinicians receive structured educational content on the topic. Leaders could consider integrating formal foundational AI teachings into medical school curricula and, then in later years, providing a chance to discuss its implications and applications in health care. These sessions could also be incorporated into postgraduate education and into continuing medical education sessions offered by workplaces.

Acknowledgments

The authors declare the use of generative artificial intelligence (GAI) in the research and writing process. According to the Generative Artificial Intelligence Delegation Taxonomy (2025), the following tasks were delegated to GAI tools under full human supervision: proofreading; editing; adapting and adjusting emotional tone; and reformatting sentence structure in specific, limited portions of text. On occasion, ChatGPT was used as a search engine to find links to relevant references, with authors further reviewing and ensuring the accuracy of these references. The GAI tool used was ChatGPT (version 5.1; OpenAI). Responsibility for the final manuscript lies entirely with the authors. GAI tools are not listed as authors and do not bear responsibility for the final outcomes.

Funding

No external financial support or grants were received from any public, commercial, not-for-profit, or other entity for any part of this work.

Conflicts of Interest

None declared.

References

1. What is artificial intelligence (AI)? IBM. URL: <https://www.ibm.com/think/topics/artificial-intelligence> [Accessed 2025-06-28]
2. Shortliffe EH. A rule-based computer program for advising physicians regarding antimicrobial therapy selection. Presented at: ACM '74: Proceedings of the 1974 Annual ACM Conference; Jan 1, 1974; San Diego, CA. [doi: [10.1145/1408800.1408906](https://doi.org/10.1145/1408800.1408906)]
3. Campbell M. Knowledge discovery in deep blue. *Commun ACM*. Nov 1999;42(11):65-67. [doi: [10.1145/319382.319396](https://doi.org/10.1145/319382.319396)]
4. Lee SB. Development of a chest X-ray machine learning convolutional neural network model on a budget and using artificial intelligence explainability techniques to analyze patterns of machine learning inference. *JAMIA Open*. Jul 2024;7(2):ooae035. [doi: [10.1093/jamiaopen/ooae035](https://doi.org/10.1093/jamiaopen/ooae035)] [Medline: [38699648](https://pubmed.ncbi.nlm.nih.gov/38699648/)]
5. Lee SB. Gradual poisoning of a chest x-ray convolutional neural network with an adversarial attack and AI explainability methods. *Sci Rep*. Jul 1, 2025;15(1):21779. [doi: [10.1038/s41598-025-02294-3](https://doi.org/10.1038/s41598-025-02294-3)] [Medline: [40593872](https://pubmed.ncbi.nlm.nih.gov/40593872/)]
6. Gichoya JW, Banerjee I, Bhimireddy AR, et al. AI recognition of patient race in medical imaging: a modelling study. *Lancet Digit Health*. Jun 2022;4(6):e406-e414. [doi: [10.1016/S2589-7500\(22\)00063-2](https://doi.org/10.1016/S2589-7500(22)00063-2)] [Medline: [35568690](https://pubmed.ncbi.nlm.nih.gov/35568690/)]
7. Benito-León J, Del Castillo MD, Estirado A, Ghosh R, Dubey S, Serrano JI. Using unsupervised machine learning to identify age- and sex-independent severity subgroups among patients with COVID-19: observational longitudinal study. *J Med Internet Res*. May 27, 2021;23(5):e25988. [doi: [10.2196/25988](https://doi.org/10.2196/25988)] [Medline: [33872186](https://pubmed.ncbi.nlm.nih.gov/33872186/)]
8. Nalinthasnai N, Thammasudjarit R, Tassaneyasin T, et al. Unsupervised machine learning clustering approach for hospitalized COVID-19 pneumonia patients. *BMC Pulm Med*. Feb 8, 2025;25(1):70. [doi: [10.1186/s12890-025-03536-w](https://doi.org/10.1186/s12890-025-03536-w)] [Medline: [39923003](https://pubmed.ncbi.nlm.nih.gov/39923003/)]
9. Rosenblatt F. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev*. Nov 1958;65(6):386-408. [doi: [10.1037/h0042519](https://doi.org/10.1037/h0042519)] [Medline: [13602029](https://pubmed.ncbi.nlm.nih.gov/13602029/)]

10. McCulloch WS, Pitts W. A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys*. Dec 1943;5(4):115-133. [doi: [10.1007/BF02478259](https://doi.org/10.1007/BF02478259)]
11. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature*. Oct 1986;323(6088):533-536. [doi: [10.1038/323533a0](https://doi.org/10.1038/323533a0)]
12. Lippmann R. An introduction to computing with neural nets. *IEEE ASSP Mag*. 1987;4(2):4-22. [doi: [10.1109/MASSP.1987.1165576](https://doi.org/10.1109/MASSP.1987.1165576)]
13. Ackley DH, Hinton GE, Sejnowski TJ. A learning algorithm for Boltzmann machines. *Cogn Sci*. Mar 1985;9(1):147-169. [doi: [10.1016/S0364-0213\(85\)80012-4](https://doi.org/10.1016/S0364-0213(85)80012-4)]
14. Li M, Bi Z, Wang T, et al. Deep learning and machine learning with GPGPU and CUDA: unlocking the power of parallel computing. *arXiv*. Preprint posted online on Oct 8, 2024. [doi: [10.48550/arXiv.2410.05686](https://doi.org/10.48550/arXiv.2410.05686)]
15. Liu BD, Meng J, Xie WY, Shao S, Li Y, Wang Y. Weighted spatial pyramid matching collaborative representation for remote-sensing-image scene classification. *Remote Sens (Basel)*. 2019;11(5):518. [doi: [10.3390/rs11050518](https://doi.org/10.3390/rs11050518)]
16. Alarcon N. OpenAI presents GPT-3, a 175 billion parameters language model. *NVIDIA Developer*. Jul 7, 2020. URL: <https://developer.nvidia.com/blog/openai-presents-gpt-3-a-175-billion-parameters-language-model/> [Accessed 2026-01-17]
17. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. Presented at: Proceedings of the 31st International Conference on Neural Information Processing Systems; Dec 4-9, 2017; Long Beach, CA. [doi: [10.5555/3295222.3295230](https://doi.org/10.5555/3295222.3295230)]
18. Dowie T. Exploring the diagnostic capability of artificial intelligence in dermatology for darker skin tones: a narrative review. *Cureus*. Oct 2025;17(10):e94909. [doi: [10.7759/cureus.94909](https://doi.org/10.7759/cureus.94909)] [Medline: [41262830](https://pubmed.ncbi.nlm.nih.gov/41262830/)]
19. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. Oct 25, 2019;366(6464):447-453. [doi: [10.1126/science.aax2342](https://doi.org/10.1126/science.aax2342)] [Medline: [31649194](https://pubmed.ncbi.nlm.nih.gov/31649194/)]
20. An approach to technical AGI safety and security. *Medium*. 2025. URL: <https://deepmindsafetyresearch.medium.com/an-approach-to-technical-agi-safety-and-security-25928819fbc6> [Accessed 2025-17]
21. Norgeot B, Quer G, Beaulieu-Jones BK, et al. Minimum information about clinical artificial intelligence modeling: the MI-CLAIM checklist. *Nat Med*. Sep 2020;26(9):1320-1324. [doi: [10.1038/s41591-020-1041-y](https://doi.org/10.1038/s41591-020-1041-y)] [Medline: [32908275](https://pubmed.ncbi.nlm.nih.gov/32908275/)]
22. Predetermined change control plans for machine learning-enabled medical devices: guiding principles. *U.S. Food & Drug Administration*. 2025. URL: <https://www.fda.gov/medical-devices/software-medical-device-samd/predetermined-change-control-plans-machine-learning-enabled-medical-devices-guiding-principles> [Accessed 2026-02-25]
23. Transparency for machine learning-enabled medical devices: guiding principles. *U.S. Food & Drug Administration*. Jun 13, 2024. URL: <https://www.fda.gov/medical-devices/software-medical-device-samd/transparency-machine-learning-enabled-medical-devices-guiding-principles> [Accessed 2026-02-25]
24. Good machine learning practice for medical device development: guiding principles. *U.S. Food & Drug Administration*. 2025. URL: <https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles> [Accessed 2026-02-25]
25. Engineering fundamentals checklist. *Microsoft Open Source*. URL: <https://microsoft.github.io/code-with-engineering-playbook/engineering-fundamentals-checklist/> [Accessed 2026-01-17]
26. Koh RGL, Khan MA, Rashidani S, et al. Check it before you wreck it: a guide to STAR-ML for screening machine learning reporting in research. *IEEE Access*. 2023;11:101567-101579. [doi: [10.1109/ACCESS.2023.3316019](https://doi.org/10.1109/ACCESS.2023.3316019)]
27. Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. Aug 2021;596(7873):583-589. [doi: [10.1038/s41586-021-03819-2](https://doi.org/10.1038/s41586-021-03819-2)] [Medline: [34265844](https://pubmed.ncbi.nlm.nih.gov/34265844/)]
28. Sloss EA, Owoyemi A, Mishra AK, et al. Development of the TrendBurden survey: assessing perceived documentation burden among health professionals in the United States. *Appl Clin Inform*. May 2025;16(3):662-675. [doi: [10.1055/a-2562-0910](https://doi.org/10.1055/a-2562-0910)] [Medline: [40669861](https://pubmed.ncbi.nlm.nih.gov/40669861/)]
29. Physician administrative burden survey –final report. *Doctors Nova Scotia*; Sep 2020. URL: <https://doctorsns.com/sites/default/files/2020-11/admin-burden-survey-results.pdf> [Accessed 2026-02-25]
30. Joint Task Force to reduce administrative burdens on physicians. *Doctors Manitoba*. May 30, 2023. URL: <https://assets.doctorsmanitoba.ca/documents/Admin-Burden-Progress-Report-May-30.pdf> [Accessed 2026-02-25]
31. Acallar LJ. AI medical scribes: everything you need to know. *Heidi*. 2026. URL: https://www.heidihealth.com/blog/ai-medical-scribe?utm_source=chatgpt.com [Accessed 2026-02-25]
32. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. Presented at: 31st Conference on Neural Information Processing Systems (NIPS 2017); Dec 4-7, 2017; Long Beach, CA. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf [Accessed 2026-02-25]
33. AI for clinicians. *Epic*. URL: <https://www.epic.com/software/ai-clinicians/> [Accessed 2026-02-25]

34. AI receptionist answers the call for busy medical clinics. Hamilton Health Sciences. URL: www.hamiltonhealthsciences.ca/share/patients-test-ai-receptionist [Accessed 2026-03-20]
35. Open Evidence. URL: www.openevidence.com [Accessed 2026-03-20]
36. AI for scientific research. Elicit. URL: <https://elicit.com> [Accessed 2026-02-25]
37. Perplexity. URL: <https://www.perplexity.ai> [Accessed 2026-02-25]
38. Google DeepMind. MedGemma. URL: <https://deepmind.google/models/gemma/medgemma> [Accessed 2026-02-25]
39. Lawrence R, Dodsworth E, Massou E, et al. Artificial intelligence for diagnostics in radiology practice: a rapid systematic scoping review. *EClinicalMedicine*. May 12, 2025;83:103228. [doi: [10.1016/j.eclinm.2025.103228](https://doi.org/10.1016/j.eclinm.2025.103228)] [Medline: [40474995](https://pubmed.ncbi.nlm.nih.gov/40474995/)]
40. Cazzato G, Rongioletti F. Artificial intelligence in dermatopathology: updates, strengths, and challenges. *Clin Dermatol*. 2024;42(5):437-442. [doi: [10.1016/j.clindermatol.2024.06.010](https://doi.org/10.1016/j.clindermatol.2024.06.010)] [Medline: [38909860](https://pubmed.ncbi.nlm.nih.gov/38909860/)]
41. Farzaneh H, Junn J, Chaibi Y, et al. Deep learning-based algorithm for automatic detection of incidental pulmonary embolism on contrast-enhanced CT: a multicenter multivendor study. *Radiol Adv*. Jun 23, 2025;2(4):umaf021. [doi: [10.1093/radadv/umaf021](https://doi.org/10.1093/radadv/umaf021)] [Medline: [41058960](https://pubmed.ncbi.nlm.nih.gov/41058960/)]
42. Aref-Eshghi E, Abadi AB, Farhadieh ME, et al. DNA methylation and machine learning: challenges and perspective toward enhanced clinical diagnostics. *Clin Epigenetics*. Oct 10, 2025;17(1):170. [doi: [10.1186/s13148-025-01967-0](https://doi.org/10.1186/s13148-025-01967-0)] [Medline: [41074112](https://pubmed.ncbi.nlm.nih.gov/41074112/)]
43. Chang WS, Hsiao KY, Lin LY, Chen M, Shia BC, Lin CY. Machine learning models for predicting in-hospital cardiac arrest: a comparative analysis with logistic regression. *Int J Gen Med*. 2025;18:6341-6352. [doi: [10.2147/IJGM.S569559](https://doi.org/10.2147/IJGM.S569559)] [Medline: [41141891](https://pubmed.ncbi.nlm.nih.gov/41141891/)]
44. Drysch M, Reinkemeier F, Pusch F, et al. Streamlined machine learning model for early sepsis risk prediction in burn patients. *NPJ Digit Med*. Oct 21, 2025;8(1):621. [doi: [10.1038/s41746-025-02078-z](https://doi.org/10.1038/s41746-025-02078-z)] [Medline: [41120704](https://pubmed.ncbi.nlm.nih.gov/41120704/)]
45. Moss L, Corsar D, Shaw M, Piper I, Hawthorne C. Demystifying the black box: the importance of interpretability of predictive models in neurocritical care. *Neurocrit Care*. Aug 2022;37(Suppl 2):185-191. [doi: [10.1007/s12028-022-01504-4](https://doi.org/10.1007/s12028-022-01504-4)] [Medline: [35523917](https://pubmed.ncbi.nlm.nih.gov/35523917/)]
46. Kumar AK, Ali Y, Kumar RR, Assaf MH, Ilyas S. Artificial intelligent and internet of things framework for sustainable hazardous waste management in hospitals. *Waste Manag*. Jul 15, 2025;203:114816. [doi: [10.1016/j.wasman.2025.114816](https://doi.org/10.1016/j.wasman.2025.114816)] [Medline: [40311410](https://pubmed.ncbi.nlm.nih.gov/40311410/)]
47. Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization. 2021. URL: <https://www.who.int/publications/i/item/9789240029200> [Accessed 2026-02-25]
48. The medico-legal lens on AI use by Canadian physicians. Canadian Medical Protective Association. Sep 2024. URL: <https://www.cmpa-acpm.ca/en/research-policy/public-policy/the-medico-legal-lens-on-ai-use-by-canadian-physicians> [Accessed 2026-02-25]
49. Solomonides AE, Koski E, Atabaki SM, et al. Defining AMIA's artificial intelligence principles. *J Am Med Inform Assoc*. Mar 15, 2022;29(4):585-591. [doi: [10.1093/jamia/ocac006](https://doi.org/10.1093/jamia/ocac006)] [Medline: [32722749](https://pubmed.ncbi.nlm.nih.gov/32722749/)]
50. Artificial intelligence (AI) in medical practice. College of Physicians and Surgeons of Saskatchewan. URL: <https://www.cps.sk.ca/imis/ContentBuddyDownload.aspx?DocumentVersionKey=879928d2-6caa-4894-a866-86fba85ce57> [Accessed 2026-01-19]
51. Livieri G, Mangina E, Protopapadakis ED, Panayiotou AG. The gaps and challenges in digital health technology use as perceived by patients: a scoping review and narrative meta-synthesis. *Front Digit Health*. Mar 27, 2025;7:1474956. [doi: [10.3389/fdgth.2025.1474956](https://doi.org/10.3389/fdgth.2025.1474956)] [Medline: [40212901](https://pubmed.ncbi.nlm.nih.gov/40212901/)]
52. Morris MR, Sohl-Dickstein J, Fiedel N, et al. Levels of AGI for operationalizing progress on the path to AGI. arXiv; Preprint posted online on Nov 4, 2023. [doi: [10.48550/arXiv.2311.02462](https://doi.org/10.48550/arXiv.2311.02462)]
53. Xu Y, Liu X, Cao X, et al. Artificial intelligence: a powerful paradigm for scientific research. *Innovation (Camb)*. Oct 28, 2021;2(4):100179. [doi: [10.1016/j.xinn.2021.100179](https://doi.org/10.1016/j.xinn.2021.100179)] [Medline: [34877560](https://pubmed.ncbi.nlm.nih.gov/34877560/)]

Abbreviations

- AI:** artificial intelligence
- DL:** deep learning
- eGFR:** estimated glomerular filtration rate
- LLM:** large language model
- ML:** machine learning

Edited by Amy Schwartz, Matthew Balcarras; peer-reviewed by Fumitoshi Fukuzawa, Rahul R Kumar; submitted 04.Oct.2025; final revised version received 05.Feb.2026; accepted 06.Feb.2026; published 30.Mar.2026

Please cite as:

Lee SB, Carter AB, Haider MH, Ko SB

An Introduction to AI for Clinicians: Tutorial

Interact J Med Res 2026;15:e85266

URL: <https://www.i-jmr.org/2026/1/e85266>

doi: [10.2196/85266](https://doi.org/10.2196/85266)

© Stephen B Lee, Alexis B Carter, Muhammad Hamis Haider, Seok-Bum Ko. Originally published in the Interactive Journal of Medical Research (<https://www.i-jmr.org/>), 30.Mar.2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in the Interactive Journal of Medical Research, is properly cited. The complete bibliographic information, a link to the original publication on <https://www.i-jmr.org/>, as well as this copyright and license information must be included.