January 2020

The Impact of Artificial Intelligence on Clinical Judgment:

A Briefing Document

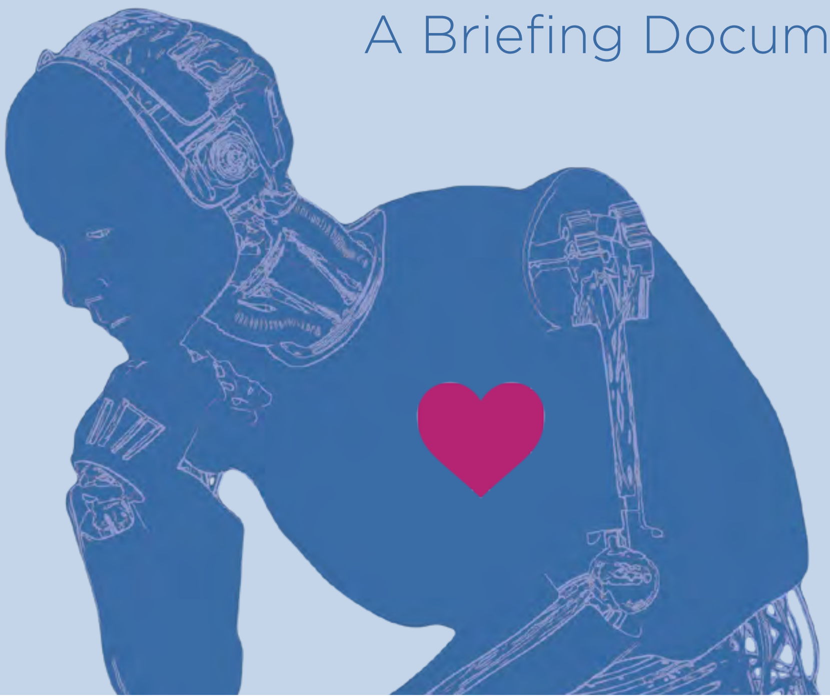

Prepared by **Benjamin Chin-Yee** and **Ross Upshur**

AMS healthcare
Transforming Education and Practice

The Impact of Artificial Intelligence on Clinical Judgment: A Briefing Document

Prepared by:

Benjamin Chin-Yee MD, MA

· Resident Physician, Department of Medicine, University of Toronto ·

and

Ross Upshur BA(Hons.), MA, MD, MSc, CCFP, FRCPC

· Dalla Lana Chair in Clinical Public Health ·
· Head, Division of Clinical Public Health, Dalla Lana School of Public Health ·
· Scientific Director, Bridgepoint Collaboratory for Research and Innovation ·
· Associate Director, Lunenfeld Tanenbaum Research Institute, Sinai Health Systems ·
· Professor, Department of Family and Community Medicine and
Dalla Lana School of Public Health, University of Toronto ·

Briefing document sponsored by AMS healthcare Transforming Education and Practice

This briefing document is one of a series of reports sponsored by AMS Healthcare. The goal of these reports is to inform health professionals, educators, and leaders about the opportunities and challenges of emerging health technologies.

AMS Healthcare is a Canadian charitable organization with an impressive history as a catalyst for change in healthcare. AMS strives to improve the healthcare of all Canadians by innovating education and practice, championing the history of medicine and healthcare, supporting leadership development, and advancing research in the health humanities.

The AMS Compassionate Care in a Technological World Initiative focuses on: (1) promoting the education and practice of compassionate care; (2) fostering new delivery models of compassionate care and (3) facilitating the leadership needed to realize the promise of technology, while safeguarding humanistic care in our rapidly evolving healthcare system.

How to cite:

Chin-Yee B, Upshur R. (2019). The Impact of Artificial Intelligence on Clinical Judgment: A Briefing Document. Toronto, ON, AMS Healthcare.

Contents

Executive Summary

- Artificial Intelligence (AI) is finding increasing applications across the healthcare system. Some predict that AI will ultimately replace the work of physicians. Others argue that human reasoning and presence will remain indispensable. In this report, we relate the current and potential applications of AI to the practice of clinical judgment. Moving beyond unhelpful dichotomies, we explain the specific ways that AI is likely to support, but not to replace, the judgment of physicians.
- To date, much of the focus has been on AI's ability to enhance predictive accuracy or augment specific diagnostic procedures, with less attention devoted to AI's potential impact on core aspects of the medical profession. Clinical judgment encapsulates the fundamental day-to-day practice of our profession and is also a major focus of undergraduate and postgraduate medical education. There is a growing need to explore how AI might interact with physicians' clinical judgment—how such technologies could support or alter clinical reasoning processes, reshape the doctor-patient relationship, and change experiential learning for future generations of physicians.
- In order to understand the potential impact of AI, it is necessary to examine the nature of clinical judgment. Clinical judgment is pluralistic in nature and requires integration of a diversity of approaches, from statistical methods to narrative frameworks. Sound clinical judgment requires tailoring the approach to what is demanded in a particular clinical circumstance. This process requires a capacity for flexible and contextual reasoning, termed "practical wisdom." The rise of AI occasions reflection on how technology might alter learners' abilities to acquire such wisdom.
- AI and machine learning are continuous with older statistical techniques, and, similar to these methods, carry a number of different strengths and limitations. Current AI applications are mostly found in diagnostic specialties and are directed at fairly narrow tasks in a manner similar to other widely used risk stratification tools or diagnostic algorithms. Although some envision a future where AI might perform more comprehensive clinical tasks, existing technologies and clinical workflows limit the likelihood that AI can effectively replace human physicians.
- Human physicians will remain indispensable because of the pluralistic nature of clinical judgment. Complex clinical tasks, from explaining a diagnosis to communicating prognosis to discussing treatment goals, require physicians to employ multiple reasoning strategies. Physicians must be equipped to apply quantitative tools alongside interpretive methods, exercising practical wisdom to determine which approach is best suited to the particular context.
- AI might offer opportunities to "humanize" medicine by helping to offset many routine tasks performed by physicians, allowing greater attention to be devoted to the irreducibly human elements of clinical medicine. To achieve this end, however, AI applications will have to be carefully designed, with sensitivity to both clinician and patient experience.
- The rise of AI and machine learning in medicine also creates opportunities to realign the goals of undergraduate and postgraduate medical education with future healthcare needs. Teaching on AI can be introduced in curricula alongside other statistical methods with the aim of developing skills in critical appraisal to help trainees assess the validity of AI tools. Programs should also foster an awareness of the social and ethical issues that arise with the use of AI in healthcare.
- AI technologies may lead to higher value being placed on uniquely human skills and capabilities, such as moral reasoning, empathy, and altruism. As a result, programs should focus on developing the humanistic competencies that underwrite sound clinical judgment. These humanistic competencies can be supported by incorporating concepts from the social sciences and humanities into core curricula at both undergraduate and postgraduate levels.

- Medical education in the era of AI should aim to support trainees in becoming technically competent, empathetic physicians who are poised to act as leaders in healthcare, finding new ways to realize the benefits of technology alongside the human capacities that enable more compassionate, patient-centered care.

1. Introduction

Artificial intelligence (AI) technologies promise to revolutionize healthcare through a range of applications that will impact how physicians make diagnoses, determine prognoses, and prescribe treatments. AI has generated interest across the healthcare system with significant implications for clinical research and quality improvement, health policy and public health, and health professions education. This interest is reflected by a growing list of publications on AI and healthcare in the form of academic articles,[1] health policy documents,[2] professional society statements,[3] and popular media coverage.

In healthcare, the term "AI" refers to a diverse set of technologies at various stages in development. The use of AI in healthcare is not a new idea,[4] and current applications are continuous with a longer history of attempts to integrate computer technology into clinical care—a movement that has been met with both enthusiasm and criticism since the 1990s. More recently, however, AI technologies have gained considerable momentum with advances in "deep learning" and the growth of "big data" to fuel machine learning algorithms.[1,5,6] These developments have stimulated debates about the impact of AI in healthcare. Some proponents claim that AI could potentially offset and even replace clinicians, while critics caution of the potential negative consequences for patient care.

This briefing document considers how AI technologies are likely to affect physicians' clinical judgment. We explore several facets of this impact with specific attention to how AI may impact clinical care and the patient-physician relationship, in addition to its implications for medical education and the training of physicians in clinical environments. Importantly, this briefing aims to move beyond polarized debates between proponents and critics of AI in healthcare, seeking rather to establish meaningful points of dialogue between parties while highlighting the challenges, risks, and opportunities for use of these technologies in clinical care.

1.1. Clinical Judgment and the Medical Profession

To date, much of the focus has been on AI's ability to perform specific technical tasks, such as enhancing predictive accuracy or augmenting specific diagnostic procedures, with less attention devoted to AI's potential impact on more central aspects of the medical profession. The concept of clinical judgment is situated at the very core of our profession, describing a fundamental activity in day-to-day practice whose development is a major focus of undergraduate and postgraduate medical education.

Clinical judgment is central to a physician's professional identity. The CanMEDS framework from the Royal College of Physicians and Surgeons of Canada illustrates the centrality of clinical judgment in the medical profession.[7] Clinical judgment exists at the intersection of "Medical Expert" and other "non-Medical Expert" roles, which include "Professional," "Communicator," and "Collaborator." It captures many key CanMEDS competencies, ranging from the ability to "perform a patient-centred clinical assessment and establish a management plan," to the capacity to "demonstrate a commitment to patients by applying best practices and adhering to high ethical standards." Clinical judgment enables physicians to "recognize and respond to the complexity, uncertainty, and ambiguity inherent in medical practice." Clinical judgment's ethical dimension ties it to the core virtues of physicianhood, recognized by the Royal College as "honesty, integrity, humility, commitment, compassion, respect, altruism, respect for diversity." Clearly, clinical judgment goes beyond performance of isolated cognitive or procedural skills to encompass a more comprehensive and diverse set of technical and human capacities.

We examine this in more detail below in our discussion of different ways to approach and conceptualize clinical judgment.

The centrality of clinical judgment in the medical profession creates a need to examine its potential interactions with AI. How might AI technologies support or alter clinical reasoning processes, reshape the doctor-patient relationship, and change experiential learning for future generations of physicians?

1.2. Overview of the Briefing Document

This briefing document is organized around four central questions that help guide our discussion:

1. **What is AI in healthcare?**
2. **What is clinical judgment?**
3. **How will AI impact clinical judgment?**
4. **What challenges and opportunities does AI create for medical education?**

We address each question in turn, beginning with an introduction to key concepts and useful terminology for understanding the role of AI in healthcare.

2. What is AI in Healthcare?

AI has become a buzzword in contemporary healthcare, and hype surrounding AI can often obscure realistic assessment of its current and potential future clinical applications. Many practicing physicians may be unfamiliar with AI terminology; therefore, we will begin by defining some key terms.[8]

Artificial Intelligence is an interdisciplinary field spanning computer science, psychology, linguistics, and philosophy, among others, focused on creating computers that can perform tasks normally associated with human intelligence.

Machine Learning is a branch of AI that uses computer systems to "learn" patterns and construct algorithms from large amounts of data. Machine learning is said to differ from older, logic-based, expert systems (so-called "Good Old-Fashioned AI") in that the algorithm's "knowledge" is derived from data, rather than from pre-programmed rules.

Deep Learning is a type of machine learning that uses multi-layered **Neural Networks** to extract features from data and generate representations at increasing levels of abstraction. Neural networks are loosely modelled after the biological nervous system. They are composed of layers of interconnected nodes which, like neurons, receive inputs and generate outputs when given thresholds are reached. In a neural network, information is encoded in the connections between nodes by how inputs are differentially weighted and their relations to the outputs generated. Connections between nodes are "trained" to represent the data in "hidden layers" that encode the relationship between inputs and outputs.

Machine learning can be further subdivided into **Supervised Learning**, which trains algorithms to classify inputs based on a desired output (i.e., based on a known classification or "ground truth"), and **Unsupervised Learning**, which uses algorithms to cluster inputs based on similar features. One commonly cited example of supervised learning is a system that trained a deep neural network to identify skin cancers based on an image database labelled with known, biopsy-proven diagnoses.[9] Unsupervised machine learning has been applied to generate novel classifications of tumours from genomic data.[10] To date, most AI applications in healthcare use supervised learning algorithms, which require labelled data and pre-specified outputs.

The relationship between these AI terms is represented in Figure 1, adapted from Goodfellow et al.[11]

Deep learning has generated considerable interest in healthcare, and receives the majority of attention in recent reviews of AI and machine learning for medical audiences.[12,13] This may be in part due to high-profile examples of deep learning applications in medicine, in particular for image recognition such as the identification of skin cancer or diabetic retinopathy.[9,14] There are, however, several other forms of machine learning algorithms that have found applications in medicine.

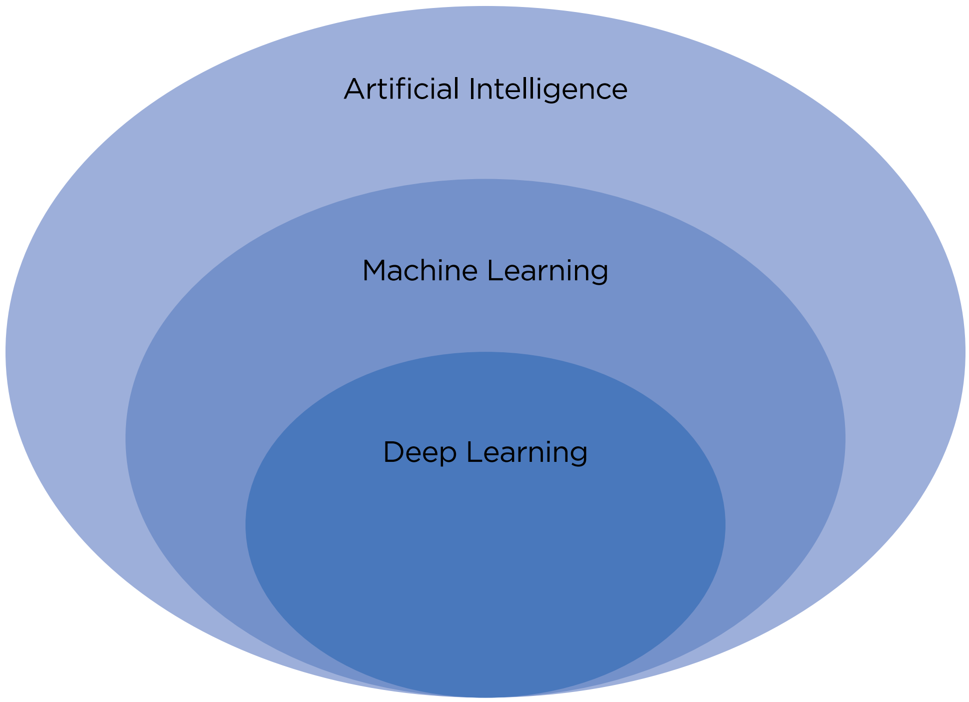

Figure 1. Schematic diagram adapted from Goodfellow et al. [11] demonstrating relationship between the terms "Artificial Intelligence," "Machine Learning" and "Deep Learning."

2.1. AI and Traditional Statistics

Despite its perceived novelty and power, machine learning is continuous with statistical methods long employed in clinical research and practice. As data scientists Andrew Beam and Isaac Kohane point out in a recent issue of JAMA[15]:

> *"Machine learning is not a magic device that can spin data into gold, though many news releases would imply that it can. Instead, it is a natural extension to traditional statistical approaches."*

Although machine learning may involve relatively less human input into how data is processed compared to traditional statistics, it is ultimately a mode of data analysis which involves performing a complex set of mathematical operations and generating inferences based on data. Machine learning may differ in its degree of automation and scale, but it does not differ in principle from other forms of mathematical analysis.

Like statistics, different methods carry advantages and disadvantages. Different machine learning algorithms offer varying degrees of intelligibility, reproducibility, and robustness. Deep learning, for example, has both strengths and limitations. It is well suited for data-rich problems, such as image classification, where there are curated, labelled datasets with known "ground truths," i.e., known values for target variables that can be used for training in supervised learning algorithms. For example, in the abovementioned skin cancer study,[9] researchers applied deep learning to a database of images pre-labelled as benign or malignant based on skin biopsies. Deep learning is often praised for its robustness and stability in the face of perturbations, which has been referred to as "graceful degradation." Graceful degradation means that a system's performance becomes progressively, but not catastrophically, worse as components are perturbed or destroyed, in contrast to older rule-based programs where removing one line of code can often cause whole systems to crash. However, studies have shown that neural networks can be sensitive to slight changes in inputs, such as image positioning or orientation, which can result in drastically different outputs and significant errors. This problem is relevant in healthcare where specific changes in inputs may arise, for example, as a result of technical factors in image acquisition. While a trained radiologist can easily recognize these changes—an alteration in image angle or rotation—this might be overlooked by an algorithm and lead to erroneous outputs. Addressing such issues is imperative before these algorithms are used in clinical care.

While certain tasks may be suitable for machine learning algorithms, other methods may be preferred in situations where interpretability and human understanding are required.[16] Selecting approaches most appropriate for the task at hand will be critical if AI is to be successfully integrated into clinical practice. We will return to this point in later sections.

2.2. Current and Potential Clinical Applications of AI

AI algorithms have attracted the most attention in diagnostic specialties, such as radiology and pathology, but have also been applied to support diagnostic procedures in a range of other fields, including dermatology, ophthalmology, gastroenterology, and cardiology. Some of these applications are summarized in recent reviews conducted by cardiologist and medical futurist Eric Topol.[13] Most of these technologies have been used in research settings and have yet to be applied in routine clinical practice. However, some applications have recently been approved by the U.S. Food and Drug Administration for clinical use.[17,18] These include a system to detect diabetic retinopathy (the Idx-DR system) and software to detect stroke on CT images (Viz.AI).

Beyond these high-profile examples, AI algorithms are already being used in a variety of auxiliary settings impacting clinical care. For example, some microbiology laboratories have integrated AI-powered image recognition software to digitally analyze bacterial growth on agar plates.[19,20] Likewise, hematology laboratories are increasingly employing digital image recognition tools that use neural networks to automate blood cell differential and morphological analysis, processes previously done manually by laboratory technologists.[21] Outside of the laboratory, some hospitals are developing AI algorithms based on local data for prediction of health outcomes and identification of patients at risk for adverse events.[22]

As these examples highlight, the application of AI technologies in medicine is already well underway. Nonetheless, AI's current applications are mostly directed at fairly narrow tasks, playing a supportive role which is often one or more steps removed from the clinical encounter. Most are designed to help support diagnostic or prognostic decision-making, and all require human interpretation for use in patient care. In this way, such algorithms are not so different from other widely used risk stratification scores and diagnostic algorithms, which have been integrated into clinical practice in many specialties and are now readily available on smartphone apps. While novel AI technologies leverage more data, employ greater statistical sophistication, and offer higher degrees of automation, in principle they serve a similar purpose to clinical tools already in existence.

Further adoption of AI technologies across health systems has been limited by a lack of prospective validation, which raises concerns about their accuracy, generalizability, and safety outside of research settings. AI diagnostic algorithms tend to perform worse when applied prospectively in comparison to their initial retrospective validation. For example, an algorithm to predict risk of in-hospital mortality developed using historical data from one centre may not perform equally well when applied at another centre, or even if applied prospectively at the same centre where it was retrospectively validated. A variety of factors may be responsible for this, ranging from bias in sample populations and idiosyncrasies in datasets to more general issues that arise when attempting to predict the future based on the past in complex systems.[23] Although these problems are not unique to AI and are well recognized in other statistical domains, they risk being overlooked in the face of hype and false assumptions that AI can "magically" transcend the limits of mathematics and traditional statistics.

Without sufficient prospective validation, it remains unlikely that standardized AI tools will become part of routine clinical care. With respect to generalizability, some AI algorithms may remain local, having been developed from local data and designed to influence local decision-making. The contextual nature of these algorithms may prevent their straightforward extrapolation to other healthcare environments. Some AI technologies have been successfully integrated into workflows of particular clinical laboratories; however, many require considerable accompanying infrastructure and support, creating challenges for transporting them to other settings. A range of factors are at play which influence how AI technologies will be integrated into clinical care moving forward.

Keeping these limitations in mind, current applications still provide the best indication of the potential future role of AI in healthcare. Although future developments in technology are notoriously difficult to predict, there will most likely be growth of narrow AI applications, such as algorithms designed to support particular diagnostic procedures or to assist in risk assessment and prognostication. Some envision a future where AI might perform a more comprehensive set of clinical tasks. For example, digital health company Babylon Health has invested in developing an AI-powered app to triage and diagnose patients.[24] The efficacy and safety of this technology and other similar tools remain to be seen.

Despite these ambitions, it is unlikely that AI will effectively replace physicians. AI will play a supportive role but will still require human interpretation for use in patient care. This is not only because of technical limitations and issues of validation and generalizability but also because of the nature of clinical medicine, where many tasks elude description in the form of a rule or algorithm. While an AI algorithm might make the correct diagnosis in a high proportion of cases, communicating that diagnosis and implementing a management plan for a particular individual, often with complex comorbidities and social care needs, requires a uniquely human skill. This human skill, which forms the core of our profession, we refer to as clinical judgment.

3. What is Clinical Judgment?

An agreed upon definition of clinical judgment remains contentious, and there are numerous perspectives on what constitutes clinical judgment articulated in the literature.[25] One definition is offered by medical ethicist Tristram Engelhardt, who defined clinical judgment as "the ability to form diagnoses, forward prognoses, and make choices of treatment which help the patient or which at least do him or her no harm".[26] In this briefing document, we employ a similar broad definition:

> *Clinical judgment refers to the range of complex reasoning tasks and actions performed by clinicians in the context of offering diagnosis, therapeutic options, and prognosis to patients regarding their health and illness.*

"Clinical" relates to the world of practice: clinical judgment is based in encounters with patients, families, and caregivers. "Judgment" refers to a heterogeneous collection of tasks, from formulating a differential diagnosis to initiating a management plan to communicating a prognosis. Clinical judgment necessarily considers multiple sources of information as well as multiple constraints on a decision. Many healthcare professionals (such as nurses, occupational therapists, and pharmacists) also engage in clinical judgment; however, here we focus on clinical judgment as performed by physicians. We use the term clinical judgment interchangeably with clinical reasoning but prefer the former, which also captures noncognitive and value-laden aspects of judgment that we discuss below.

3.1. Approaches to Clinical Judgment

Clinical judgment has been studied from multiple perspectives. Some have applied insights from cognitive science and psychology to examine clinical reasoning in terms of distinct cognitive processes, highlighting how common "heuristics" and "biases" influence medical decisions.[27-29] This approach has been influential in medical education and clinical practice, in particular for thinking about medical errors. Others have examined tacit knowledge, arguing that intuition constitutes a critical and indispensable component of clinical judgment.[30,31]

Another approach emphasizes the ethical, value-laden dimensions of clinical judgment. This view highlights the clinician's role as moral reasoner, with clinical judgment requiring the cultivation of virtues and practical wisdom. In her book How Doctors Think,[32] humanities scholar and medical educator Kathryn Montgomery defines **Practical Wisdom** or **Phronesis** as:

> *"The flexible, interpretive capacity that enables moral reasoners to determine the best action to take when knowledge depends on the circumstance."*

Practical wisdom, or phronesis, is the essential virtue which "enables physicians to fit their knowledge and experience to the circumstances of each patient." Recognizing practical wisdom as a core component of clinical judgment has implications for how we conceive of the role of AI in medicine:

> *"If medicine were a science in the old-fashioned positivist sense, its laws could be programmed, and diagnosis could be determined, and choice of treatment decided entirely by computer. There would be no need for physicians. But even if computer programs ... worked most of the time, they would still be an inadequate substitute for clinical attention. The need for human contact by both parties to the patient-physician encounter goes well beyond the patient's need for reassurance and support. Clinicians need to examine the patient for themselves ... What experienced clinicians possess ... is an immense and well-sorted catalogue of clinical cases and the clinical judgment to know how to use it, and that store of knowledge is activated by seeing, touching, and questioning the patient."*[32]

Related to these approaches are narrative accounts, which explore the interpretive processes by which physicians and patients engage in shared decision-making. These accounts acknowledge the subjective experience of illness and its entanglement within "history, culture and life-meaning".[33] Sound clinical judgment is underwritten by "narrative skills of recognizing, absorbing, interpreting, and being moved by the stories of illness."

Lastly, one of the most influential approaches to clinical judgment was popularized by the Evidence-Based Medicine (EBM) movement in the 1990s. This approach championed statistical methods for clinical reasoning, laying down algorithmic procedures, which are exemplified by the multitude of risk scores and clinical prediction rules that emerged. Clinical decision-making algorithms, such as the Wells Score for prediction of Pulmonary Embolism risk, have had a significant impact in medical practice. Risk scores have gone a long way in stratifying patients, helping determine appropriate diagnostic testing, and guiding treatment. Systematic reviews have demonstrated how standardized decision support algorithms can improve diagnosis and management. As a result, these tools have been incorporated into clinical practice guidelines and have become standards for teaching clinical diagnosis in undergraduate and postgraduate medical curricula.

3.2. Statistical versus Interpretive Frameworks

How do these different approaches to clinical judgment hang together? A broad distinction can be drawn between "statistical" and "interpretive" approaches to clinical reasoning (Figure 2). A similar distinction can be found in the work of psychologist Paul Meehl, whose 1954 book Clinical versus Statistical Prediction contrasted "clinical" judgment with statistical or "actuarial" judgment, the former being described as "subjective" and the latter as "objective." Meehl argued that reasoning aided by some form of mathematical model, usually based in statistics or probability, was superior to unaided reasoning, which was subject to numerous biases, a finding echoed by research in cognitive psychology mentioned above.

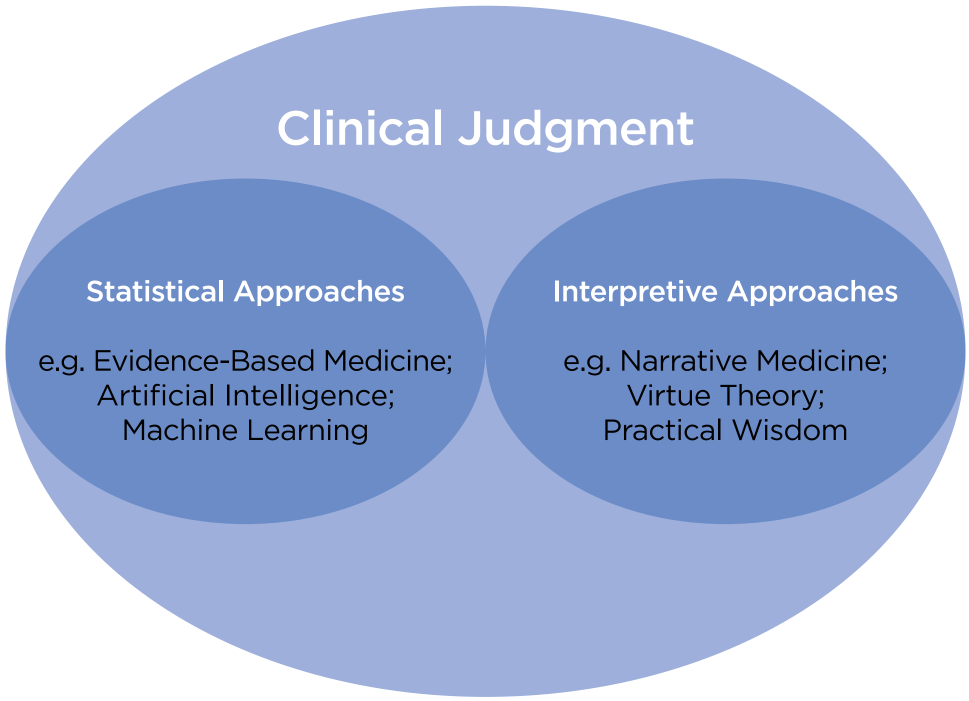

Figure 2. Schematic diagram illustrating different approaches to clinical judgment and distinguishing between statistical and interpretive frameworks.

Much of the thrust of EBM has been to support statistical approaches to clinical judgment. Importantly, many EBM algorithms do leave space for clinical intuition, not only in how rules are applied in practice but also within the scores themselves. For example, the Wells Score includes "subjective" criteria, "pulmonary embolism is as likely as or more likely than an alternative diagnosis," which allows additional points to be allocated based on clinical experience or gestalt to place patients in a higher risk category. This need for clinical expertise to inform clinical reasoning was recognized early on by EBM proponents. As EBM pioneer David Sackett and colleagues stated in their 1996 BMJ article[34]:

> *"External clinical evidence can inform, but can never replace, individual clinical expertise, and it is this expertise that decides whether the external evidence applies to the individual patient at all and, if so, how it should be integrated into a clinical decision."*

EBM's algorithmic focus does not obviate the need for interpretation by physicians.[35] This need establishes a role for narrative skills and practical wisdom in accounts of clinical judgment. Such accounts can help physicians judge the best decision within a given circumstance and help to guide ethical action while tailoring care to patient needs and experiences.

4. How Will AI Impact Clinical Judgment?

Where do AI and machine learning fit in with these approaches to clinical judgment? Given their reliance on vast amounts of digitized data and statistical analysis, AI and machine learning can be seen as extensions of statistical approaches to clinical reasoning.

Machine learning, however, differs from EBM algorithms in that rather than data being inputted into predefined algorithms, AI can derive the algorithms from data sets. A major impetus behind EBM's quantitative tools was to render clinical reasoning processes explicit—to make clear the inputs into decision-making and show how they impact the action advised. In this way, many AI tools depart from older clinical prediction rules and risk scores. With such tools, especially those using deep learning neural networks, it is not always possible for a physician to untangle how different inputs lead to changes in outputs to shape clinical decision-making.

4.1. Clinical Judgment, AI and Black Boxes

The potential unintelligibility of AI is an issue often raised in the literature, with critics claiming that an AI model constructs a "black box" that prevents clinicians from accessing and understanding the reasons behind a particular decision. As leading computer scientist and deep learning researcher Geoffrey Hinton put it in a recent issue of JAMA[5]:

> *"Understandably, clinicians, scientists, patients, and regulators would all prefer to have a simple explanation for how a neural net arrives at its classification of a particular case. In the example of predicting whether a patient has a disease, they would like to know what hidden factors the network is using. However, when a deep neural network is trained to make predictions on a big data set, it typically uses its layers of learned, nonlinear features to model a huge number of complicated but weak regularities in the data. It is generally infeasible to interpret these features because their meaning depends on complex interactions with uninterpreted features in other layers."*

Some have argued that concerns over black boxes are exaggerated, given that medicine is fraught with uncertainties and many decisions made by physicians rely on features that cannot be explicitly stated or explained, which remain outside the realm of medical evidence.[36] In other words, AI may just be replacing one black box, call it clinical intuition, with another.

We should not, however, underestimate the need for explanation and justification in medicine where decisions impact patient care. The need to provide reasons for clinical decisions, or to ground interpretations of a patient's illness, is a core ethical responsibility of clinicians. Not everyone is convinced that the analogy between the unintelligibility of machines and human reasoning holds, and that both situations are ethically equivalent. Ultimately human decision makers are accountable for their actions. The same—at least presently—cannot be said for AI, although the problem of machine accountability remains the subject of active debate. If AI applications are to become widespread in medicine, further critical discussions surrounding these issues will be needed.

4.2. Interpretation and Reciprocity in AI

Despite the important differences between AI and traditional EBM-inspired approaches, situating AI alongside other statistical methods helps us better delimit its potential role in clinical judgment. The influence of AI may be understood as continuous with other mathematical approaches to clinical reasoning. Technologies using AI may help to assess patients' risk, determine when further diagnostic work-up is appropriate, and suggest treatment options. But like other quantitative tools, AI algorithms will always require some element of human interpretation. The extent of interpretation required will vary with each specific application and how it relates to patient care.

Some AI applications, such as the aforementioned laboratory technologies, may require less interpretation on the part of the clinician. In the laboratory, AI technologies are replacing procedures that do not usually involve physician interpretation, which in part explains why the integration of AI has been less contentious in these settings. For example, the use of validated machine learning algorithms to identify positive bacterial cultures is seen as unproblematic by most physicians. (Indeed, many physicians are likely unaware of the specific technologies being used.) Physicians take this machine-generated output, indicating a “positive” or “negative” culture, at face value. Certainly, this output must be then interpreted within the clinical context, for example, to guide antimicrobial selection. Here, clinical judgment weighs multiple factors—from patient-level factors (such as allergies, immune status, or concomitant organ dysfunction) to broader systems-level considerations (such as resource utilization or risk of antimicrobial resistance).

AI technologies may more easily replace procedures for which minimal physician interpretation is needed. However, minimal interpretation does not mean that interpretation is removed altogether. Clinical practice requires a reciprocal exchange between physicians and technology. This becomes apparent when the reliability of an output is questioned: for instance, when a particular result faces scrutiny in light of clinical context. Most physicians take the automated complete blood count at face value and some centres now use a neural network-generated white blood cell differential. However, when the results are discrepant or unexpected, manual count and morphological analysis are requested and performed under the microscope through the trained eye of a technologist or hematopathologist.

Likewise, many clinicians take a hands-off approach when it comes to imaging results: when a radiologist’s report accords with their clinical expectation then they need look no further. However, when a report details an unexpected finding, or one that will significantly change management, this occasions a dialogue between clinician and radiologist. Such a dialogue is often mutually illuminating: the clinician provides additional clinical context, allowing the radiologist to reinterpret a finding and narrow the differential diagnosis. Likewise, radiologists and pathologists, when confronted with a puzzling image or finding under the microscope, will routinely review the clinical history and call the most responsible physician for additional clinical information to aid in their interpretation. Such exchanges underscore the need for explanation and justification when making decisions that impact patient care. The practice of medicine involves continuous interplay between parties—physician-patient, physician-physician, physician-technologist—all happening in real time as information and context changes, and reinterpretation is required. While some pundits argue that AI will replace radiologists and pathologists, in our view, these examples highlight how such specialists will continue to play a critical role, even in the face of advanced image recognition software.

The ongoing need for reciprocal exchange between technologies and their users in clinical decision-making raises concerns about the impact of new technologies that threaten to reduce or eliminate this reciprocity. Amidst growing technical and clinical pressures, physicians already find it challenging to engage consistently with technologies and technologists in ways that enrich interpretation and judgment. Work is needed to preserve this reciprocity in healthcare. As the renowned Canadian physicist and author Ursula Franklin warned[37]:

> *“Whenever human activities incorporate machines or rigidly prescribed procedures, the modes of human interaction change. In general, technical arrangements reduce or eliminate reciprocity. Reciprocity is some manner of interactive give and take, a genuine communication among interacting parties … Once technical divides are interposed, they allow a physical distance between the parties. The give and take—that is, the reciprocity—is distorted, reduced, or even eliminated.”*

Loss of reciprocity is not a welcome development. Reciprocity is required wherever flexible, contextual reasoning is demanded, and medicine is a paradigm example of a domain where this is the case. As Franklin points out:

> *"Reciprocity ... is situationally based. It's a response to a given situation. It is neither designed into the system nor is it predictable. Reciprocal responses may indeed alter initial assumptions. They can lead to negotiations, to give and take, to adjustment and they may result in new and unforeseen developments."*

The need for reciprocity and interpretation is most apparent when one applies mathematical tools in patient-facing, clinical settings. A putative AI algorithm may produce as outputs a number of diagnostic or treatment options, but clinical judgment involves relating these options to an individual patient by, for example, aligning them with the patient's preferences and values. This process involves a reciprocal "give and take" to negotiate which options are best suited to the individual in their life circumstance. Furthermore, counselling a patient on the different options, and having them buy in to a given management strategy, requires careful explanation and demonstration of how therapeutic goals align with their values and personal goals.

Nuanced variables such as values and goals are difficult to code as inputs into quantitative models, leaving much of medicine's rich and informative qualitative data systematically excluded from big data approaches. Instead, assessing these variables requires engagement with first-person experience, an essential part of clinical judgment emphasized by narrative approaches. Clinical judgment remains a reciprocal, interpretive practice that necessitates the integration of both quantitative and qualitative reasoning strategies.

4.3. Getting Beyond AI "versus" Clinical Judgment

Conflict between AI and clinical judgment is a recurring theme in the academic literature and popular media,[38] often leading parties to take on polarized positions. Both sides are to blame for this polarization. From proponents who make sweeping claims about far-fetched scenarios to critics who stoke fears of replacement of physicians by machines, tensions have arisen that can hinder opportunities for sober assessment of how AI will affect clinical judgment.

Rather than assuming AI inherently conflicts with clinical judgment, we should understand that AI develops a particular component judgment, namely the statistical approaches detailed above. AI applications will contribute a set of quantitative tools that exist alongside other resources and methods used in clinical decision-making. In areas where quantitative approaches predominate, such as in the laboratory, AI technologies may have a significant impact. Here AI technologies will be more readily integrated into clinical workflows, impacting patient care without appearing to encroach on physicians' exercise of judgment. In patient-facing domains, however, integration of these tools requires recognition that they serve only one aspect of clinical judgment and must work alongside interpretive approaches to allow for shared decision-making and patient-centred care.

Just as physicians can no longer rely on clinical intuition alone, relying solely on quantitative methods does not provide an adequate model of clinical judgment. We must avoid elevating AI technologies as exemplars of clinical judgment while devaluing experiential knowledge and interpretive reasoning, which will continue to play a critical role in patient care.[39,40]

On a more optimistic note, some have argued that AI may help offset many cognitive tasks performed by physicians, freeing up more time for the irreducibly "human" elements of clinical medicine. Physicians will be given more time and energy to engage in patient-centred communication and express empathy, allowing for the delivery of more compassionate care. As physician and author Abraham Verghese and colleagues put it[41]:

> *"In the same manner that automated blood pressure measurement and automated blood cell counts freed clinicians from some tasks, artificial intelligence could bring back meaning and purpose in the practice of medicine while providing new levels of efficiency and accuracy. Physicians must proactively guide, oversee, and monitor the adoption of artificial intelligence as a partner in patient care."*

We cannot assume that more compassionate care will inevitably follow from the rise of AI in medicine and the increased efficiencies that such technologies might afford. AI applications will have to be carefully designed, with sensitivity to both user and patient experience, if such ends are to be achieved. Many have chronicled the unintended negative effects of the rise of electronic medical records (EMRs) on the patient-physician relationship and physician burnout.[42]

As the case of EMRs demonstrates, more efficiency does not always lead to more compassionate care. EMRs have enhanced our ability to collate data, such that physicians now have access to a wealth of information about the patient before they even enter the room. In the future, AI tools may allow physicians to use this data to make a diagnosis and determine treatment options before even meeting the patient. Such technologies may make for more efficient consultation; however, as most experienced clinicians recognize, it is often through taking the time to gather a history for oneself that new and relevant information comes to light. This is not simply a matter of information lost or gained. More importantly, this human presence and openness to learning about the other's experience is the key to building a therapeutic relationship. Listening to the story remains the foundation of the physician-patient encounter. Therefore, while it may be perceived as inefficient by some, human interaction continues to hold vital clinical and educational value.

4.4. Example: Prediction versus Prognostication

Let us now look at an example that highlights different approaches to clinical reasoning and shows how AI might contribute in practice. Prognostication is a core part of a physician's work and clinical judgment. Many have employed statistical tools to predict clinical outcomes, and there is now a plethora of prognostic models, from disease-specific indices to more general risk scores. AI and machine learning have also been applied to problems of prognostication. For example, a study from Stanford University applied a deep learning algorithm to predict all-cause mortality within three to twelve months for admitted patients based on historical EMR data.[43] Similar initiatives are currently using machine learning to predict risk of inpatient mortality at hospitals in Canada.

Such tools, if sufficiently validated in representative patient populations, may indeed prove useful for predicting risk of death, and in some cases can be paired with a clinical intervention. For instance, the Stanford group used their deep learning algorithm to make recommendations for palliative care referral for patients at high mortality risk.[43] The medical literature has seen an increasing emphasis on prediction, evidenced by the rise of Predictive Analytics, a term borrowed from business and finance, which uses big data and machine learning to predict future events under conditions of uncertainty.[44]

This focus on prediction may lead us to forget that prediction and prognostication are not synonymous. Whereas prediction is focused on an outcome, quantifiable within error limits, prognostication is a more inclusive concept, which has its roots in ancient medicine. This point is emphasized by John Thomas and colleagues in a recent article in JAMA Internal Medicine.[45] They advocate for a return to a Hippocratic concept of prognosis, whereby prognosis plays an explanatory role in patient care. According to this view, prognostication describes the process through which the physician aids the patient in making sense of their illness trajectory within the context of their past, present and future.

> *"The Hippocratic approach makes clear the great importance of understanding the contributors to predictions and how those contributors may change over time, thus emphasizing the explanatory power of prognosis. As a synthetic approach, it calls for more than a simple prediction or estimate, and in this way, it goes beyond the use of prognostic tools and does not necessarily involve quantified estimates. It also allows for the possibility of considering a variety of outcomes in addition to mortality, tailoring discussions to the outcomes that are most pertinent to the clinical situation. Sharing this synthetic understanding with patients has the potential to accomplish more than to inspire confidence; it may also have educational value for patients in communicating the relevant health-related aspects of the past and present, and how they inform the future. It would make clear the reality that prognosis is dynamic, in that relevant factors influencing prognosis may change over time, necessitating periodic reassessments and discussions."*

This passage highlights how, even with the rise of predictive analytics and machine learning algorithms, interpretive, narrative approaches will remain indispensable to prognostication. Although AI models may accurately predict some clinical outcomes, clinical judgment requires that such measures be interpreted and explained within the context of the patient's life. Prognostication cannot be reduced to an output generated by a predictive algorithm. Clinical judgment may draw upon algorithmic tools, but it also involves more nuanced reasoning processes. These processes are perhaps best captured by the notion of practical wisdom, introduced above, which describes the interpretive capacities that allow physicians to apply knowledge and experience to determine the best action in a particular circumstance. Practical wisdom, cultivated through clinical experience, is what helps guide how physicians approach questions such as: "When should prognosis be raised in a discussion with a patient?" or "When is it appropriate to apply a quantitative tool and offer numbers? When is it not?" or "How should one communicate this information, along with its attendant uncertainties, to a patient in a way that they can understand?" Answers to these questions will not be found in the form of algorithms, yet addressing them is crucial in day-to-day clinical practice. Sound clinical judgment should leverage the available quantitative resources enabled by technological advances that have become part and parcel of modern medicine; however, it cannot neglect the equally important interpretive elements that are essential to providing compassionate patient care.

The example of prognostication shows how clinical judgment integrates a diversity of approaches—that clinical judgment is a pluralistic practice. Having a diversity of approaches is not a bad thing. Rather, this diversity acknowledges how complex clinical tasks often require multiple reasoning strategies, using quantitative tools alongside interpretive methods. Embracing a pluralistic clinical judgment avoids privileging one method over the other, recognizing that each approach can play an important role depending on the clinical scenario. In the next section, we apply these insights to explore the implications for medical education.

5. What Challenges and Opportunities Does AI Create for Medical Education?

AI creates a host of challenges and opportunities for medical education. These stem from AI's potential to reshape how certain clinical tasks are performed, which in turn could engender a shift in the skills demanded of physicians. For example, to invoke the Royal College's CanMEDS framework, particular aspects of the "Medical Expert" competencies—such as "interpreting diagnostic tests" or "determining the most appropriate procedures or therapies"—may be areas where AI technologies could make significant contributions in the future.[7] As a result, knowledge and skills that were previously seen as core professional competencies for all physicians may become less central. This, in turn, should occasion refocusing on the distinctive "human" competencies whereby physicians add value to healthcare.

5.1. Refocusing Medical Education on "Human" Competencies

Medical education plays a crucial role in reorienting physicians' training towards high value, "human" competencies in order to meet future needs amid rapid technological advancement. This line of reasoning is advanced by David Li and colleagues in their recent article "Why We Needn't Fear the Machines: Opportunities for Medicine in a Machine Learning World," published in Academic Medicine.[46] They analyze the issue through an economic lens:

> *"Economic theory tells us that the value of complements to an emerging technology like machine learning will increase while the value of substitutes will decrease. The complement of machine prediction in clinical decision making is human judgment, the ability to evaluate the benefits and costs of potential treatments based on the patient's broader context. While structured clinical metrics documented in electronic medical records are typically used to build machine learning models, unique human values and social determinants of health are difficult to model. In the future, physicians will need to combine this unstructured information with machine prediction to augment their human judgement and provide high quality patient-centered care. Human capabilities such as empathy and altruism will become more valuable, but the most valuable complement will be the one that no computer can ever replace: the human presence. Offloading highly specific routine tasks to automated technology will not make a physician's complementary skills unnecessary; it will actually increase their importance and economic value."*

Far from replacing or eliminating the need for human physicians, AI and machine learning will only increase their importance by placing higher value on uniquely human skills and capabilities—capabilities which include moral reasoning, empathy, and altruism. This aligns with other arguments that AI may help to humanize medicine.[47] However, as we discussed above, if this humanizing process is to occur, it will require deliberate action on the part of both healthcare innovators and educators. Focusing on the implications of AI for medical education, Li et al. write:

> *"It is clear that the future role of physicians will likely be redefined as certain tasks become more automated. How does the current training paradigm prepare physicians for this reality? Likely, not very well—or at least, not yet. Discussion about how the workplace will be different is now more necessary than ever. It is also clear that the value proposition of a human physician needs to be defined clearly and a new training paradigm must be developed in a deliberate manner. Medical schools must teach students cutting edge medicine without losing sight of the human aspects of medicine. Curricula should focus on two core areas: improving human judgment and delivering patient-centered care."*

This emphasis on improving human judgment is especially salient to our discussion. One concern raised earlier was that the rise of AI in medicine will place undue emphasis on quantitative approaches to clinical judgment leading to subtle erosion of its humanistic components that are essential in patient care. Proponents of AI in medicine often elevate the status of these technologies as exemplars of clinical reasoning without recognizing that this approach is one among many methods of addressing clinical questions. We have shown how this view misunderstands the nature of clinical judgment. Devaluing the uniquely human, interpretive and contextual, aspects of reasoning is the exact opposite of what is needed for the future of the medical profession.

5.2. Implications for Curricular Development

To guard against this direction, it is important that trainees learn about new technologies in a balanced manner, with attention to both their promises and their limitations. Medical curricula should introduce material on AI and machine learning alongside other statistical techniques with the aim of developing skills in critical appraisal, teaching trainees to assess the validity and applicability of AI tools in various clinical settings. In addition to offering a basic technical background to understand AI applications in medicine, curricula should foster awareness of the social and ethical issues that arise with the use of AI. To this end, a growing literature has begun to explore the social and ethical dimensions of AI in healthcare, offering critical analysis of issues ranging from data privacy and trust to equity and diversity in AI.

Critically, postgraduate and undergraduate medical curricula should continue to create space for growth of the human capacities that underwrite sound clinical judgment. As we have seen, practical wisdom is central to clinical judgment, a capacity best cultivated through experiential learning paired with reflective practice. Medical educators must be sensitive to how technology might adversely affect learners' ability to develop practical wisdom. Just as EMRs have the potential to invert the clinical encounter and risk privileging electronic data above human interaction, AI tools could lead down a similar path if applied indiscriminately and unreflectively. Education in AI should not occur at the expense of fostering narrative skills that allow learners to attend to first-person experiences that escape quantitative measurement but remain essential for patient care.

The medical humanities have gone a long way in fostering narrative skills and sensitivity to first-person experiences, and there is growing interest in introducing teaching in the humanities at both undergraduate and postgraduate levels.[48] One important initiative is the "Dialogical Learning" approach proposed by Arno Kumagai and colleagues.[49] What is original about this approach is that it occurs within and alongside clinical practice, giving pause to attend to the "dialogical moments" that make up the human-side of medicine, and using these opportunities to grow both teachers' and learners' capacities in critical reflection and empathy:

> *"From an educational point of view, what does dialogical teaching and learning look like? We would argue that it is not limited to acquisition of a new skill, such as performing a lumbar puncture or reading an abdominal ultrasound scan. It cannot be taught in a lecture or a clinical skills or motivational interviewing course, nor can it be isolated in standardized or simulated environments. To avoid excessive abstraction and to give the subject of the dialogical interactions professional and personal relevance, dialogical learning should ideally be carried out in specific clinical contexts— in clinics, wards, emergency departments, and operating rooms—during moments of moral conflict, human suffering, reflection, and wonder... These dialogical moments may consist of brief exchanges captured on the posing of a thoughtful question by an attending physician, the highlighting of an everyday event, the acknowledgment of a difficult or tragic situation—a momentary encounter that may expand in reflection and memory long after the exchange. These moments are envisioned as something different from the usual delivery of "clinical pearls" during rounds: In a sense, it is not the delivery of information that is central but the opening up of thought and perceptions through questions that prompt consideration of different perspectives, approaches, and ways of living."*

Educators must recognize how the humanities play a key role in medical training and avoid relegating humanities subjects to "companion curricula" or extracurricular domains, where they too often face neglect in favour of "core" bioscientific knowledge. Ayelet Kuper and colleagues take an important step in this direction by defining a body of social sciences and humanities knowledge that underpins important competencies for physicians.[50] Their study identified twelve interrelated themes from social sciences and humanities disciplines which serve as a basis for teaching the so-called "non-Medical Expert" CanMEDS roles and are currently being used to inform curricular development.

These examples suggest ways to teach humanistic competencies in undergraduate and postgraduate programs with a rigour that is lost when we simply classify them as "soft skills," subordinate to bioscientific competencies. To be sure, advances in bioscience and quantitative methodologies have yielded important advances in modern medicine. But there is growing recognition that these forms of knowledge alone do not capture the whole of medicine. The rise of AI creates an opportunity to refocus our attention on developing the human capacities that contribute to a flexible, pluralistic clinical judgment. Trainees educated in this manner will be equipped to practice as technically competent, empathetic physicians. They will be poised to act as leaders in healthcare, finding new ways to realize the benefits of technology to better support the human capacities that enable more compassionate, patient-centred care.

6. Conclusion

This briefing document has examined the relationship between AI and physicians' clinical judgment. We have explored questions of how current and potential future AI applications will contribute to clinical decision-making, how physicians will interact with AI in healthcare settings, and how AI will offer both challenges and opportunities for medical education.

One important lesson is the need to move beyond the tensions between clinical judgment and AI that are often emphasized in the literature. While there are valid concerns about the detrimental effects AI technologies may have on clinical practice, pitting AI against clinical judgment sets up a false division and creates unnecessary barriers to finding ways AI can be integrated alongside human capacities to benefit patient care.

No single account of clinical judgment captures the diversity and complexity of the reasoning processes required to perform the day-to-day work of physicians. AI will contribute to the arsenal of quantitative approaches used in clinical judgment that play a critical role in modern medicine. Far from replacing humanistic approaches, however, well-designed and thoughtfully applied AI technologies may have the potential to increase the value of human presence in clinical care. Trainees and physicians of the future must become leaders to ensure that AI contributes to the envisioned outcome of better, more compassionate, patient care. This briefing document serves as an introduction and launching pad for further productive dialogue among physicians, educators, and healthcare innovators.

References

1 Naylor CD. On the Prospects for a (Deep) Learning Health Care System. Journal of the American Medical Association. 2018;320(11):1099-1100

2 Stinson C. Healthy Data: Policy Solutions for Big Data and AI Innovation in Health. 2018.

3 Tang A, Tam R, Cadrin-Chênevert A, Guest W, Chong J, Barfett J, et al. Canadian Association of Radiologists White Paper on Artificial Intelligence in Radiology. Can Assoc Radiol J. 2018 May;69(2):120–35.

4 Baxt WG. Application of artificial neural networks to clinical medicine. The Lancet. 1995;346(8983):1135–8.

5 Hinton G. Deep learning—a technology with the potential to transform health care. Journal of the American Medical Association. 2018;320(11):1101-1102

6 Stead W. Clinical Implications and Challenges of Artificial Intelligence and Deep Learning. Journal of the American Medical Association. 2018;320(11):1107-1108

7 Snell L, Sherbino J, Frank J. CanMEDS 2015 Physician Competency Framework. Royal College of Physicians and Surgeons of Canada.

8 Upshur R. Artificial Intelligence, Machine Learning and the potential impacts on the practice of Family Medicine: A briefing document. 2019 Jul pp. 1–16.

9 Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017 Jan 25;542(7639):115–8.

10 Capper D, Jones DTW, Sill M, Hovestadt V, Schrimpf D, Sturm D, et al. DNA methylation-based classification of central nervous system tumours. Nature. 2018 Mar 1;555(7697):469–74.

11 Goodfellow I, Bengio Y, Courville A. Deep learning. MIT Press, 2016. www.deeplearningbook.org

12 Rajkomar A, Dean J, Kohane I. Machine Learning in Medicine. N Engl J Med. 2019 Apr 4;380(14):1347–58.

13 Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine volume 25, pages44–56(2019)

14 Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. Journal of the American Medical Association. 2016 Dec 13;316(22):2402–9.

15 Beam AL, Kohane IS. Big Data and Machine Learning in Health Care. Journal of the American Medical Association. 2018 Apr 3;319(13):1317–2.

16 Tu JV. Advantages and disadvantages of using artificial neural networks versus logistic regression for predicting medical outcomes. J Clin Epidemiol. 1996 Nov;49(11):1225–31.

17 FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems. Available from: https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye

18 FDA permits marketing of clinical decision support software for alerting providers of a potential stroke in patients. Available from: https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-clinical-decision-support-software-alerting-providers-potential-stroke

19 Faron ML, Buchan BW, Vismara C, Lacchini C, Bielli A, Gesu G, et al. Automated Scoring of Chromogenic Media for Detection of Methicillin-Resistant Staphylococcus aureus by Use of WASPLab Image Analysis Software. Richter SS, editor. Journal of Clinical Microbiology. 2016 Feb 25;54(3):620–4.

20 Faron ML, Buchan BW, Coon C, Liebregts T, van Bree A, Jansz AR, et al. Automatic Digital Analysis of Chromogenic Media for Vancomycin-Resistant-Enterococcus Screens Using Copan WASPLab. Burnham CAD, editor. Journal of Clinical Microbiology. 2016 Sep 23;54(10):2464–9.

21 Briggs C, Longair I, Slavik M, Thwaite K, Mills R, Thavaraja V, et al. Can automated blood film analysis replace the manual differential? An evaluation of the CellaVision DM96 automated image analysis system. Int J Lab Hematol. 2009 Feb;31(1):48–60.

22 Early Warning for General Internal Medicine (GIM) Patients. Available from: https://www.chartdatascience.ca/early-warning-gim

23 Chin-Yee B, Upshur R. Three Problems with Big Data and Artificial Intelligence in Medicine. Perspectives in Biology and Medicine. 2019;62(2):237–56.

24 Burki T. World Report GP at hand: a digital revolution for health care provision? The Lancet. 2019 Aug 10;394(10197):457–60.

25 Braude HD. Clinical Reasoning and Knowing. In: Marcum JA, editor. The Bloomsbury Companion to Contemporary Philosophy of Medicine. 2017.

26 Engelhardt HT, Spicker SF, Towers B, editors. Clinical Judgment: A Critical Appraisal. Vol. 6. Dordrecht: Springer Netherlands; 1979.

27 Tversky A, Kahneman D. Judgment under Uncertainty: Heuristics and Biases. Science. 1974 Sep 27;185(4157):1124–31.

28 Norman GR. The epistemology of clinical reasoning: perspectives from philosophy, psychology, and neuroscience. Academic Medicine. 2000;75:S127–33.

29 Elstein AS. Clinical Problem Solving and Decision Psychology: Comment on "The Epistemology of Clinical Reasoning." Academic Medicine. 2000 Oct 1;75(10):S134.

30 Braude HD. Intuition in Medicine. University of Chicago Press; 2012.

31 Braude HD. Clinical intuition versus statistics: different modes of tacit knowledge in clinical epidemiology and evidence-based medicine. Theor Med Bioeth. 2009 Jun 23;30(3):181–98.

32 Montgomery K. How Doctors Think: Clinical Judgment and the Practice of Medicine. Oxford University Press; 2005.

33 Charon R. Narrative Medicine: Honoring the Stories of Illness. Oxford University Press; 2006.

34 Sackett DL, Rosenberg WM, Gray JA, Haynes RB, Richardson WS. Evidence based medicine: what it is and what it isn't. BMJ. 1996 Jan 13;312(7023):71–2.

35 Djulbegovic B, Guyatt G. Evidence vs Consensus in Clinical Practice Guidelines. Journal of the American Medical Association. 2019 Jul 19:1–2.

36 London AJ. Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability. Hastings Center Report. 2019 Feb 21;49(1):15–21.

37 Franklin UM. The Real World of Technology. House of Anansi; 1999.

38 Mukherjee S. A.I. versus M.D. The New Yorker. 2017 Jul 27.

39 Chin-Yee B, Upshur R. Clinical judgement in the era of big data and predictive analytics. J Eval Clin Pract. 2018 Jun;24(3):638–45.

40 Chin-Yee B, Chin-Yee I. Big Data, Precision Medicine and Person-Centered Healthcare. European Journal for Person Centered Healthcare. 2018 Vol 6 Issue 4 pp 513-515

41 Verghese A, Shah NH, Harrington RA. What This Computer Needs Is a Physician. Journal of the American Medical Association; 2018 Jan 2;319(1):19–2.

42 Verghese A. How Tech Can Turn Doctors Into Clerical Workers. The New York Times Magazine. 2018 May 16.

43 Avati A, Jung K, Harman S, Downing L, Ng A, Shah NH. Improving palliative care with deep learning. IEEE; 2017. pp. 311–6.

44 Sniderman AD, D'Agostino RB, Pencina MJ. The Role of Physicians in the Era of Predictive Analytics. Journal of the American Medical Association; 2015 Jul 7;314(1):25–6.

45 Thomas JM, Cooney LM Jr, Fried TR. Prognosis Reconsidered in Light of Ancient Insights—From Hippocrates to Modern Medicine. JAMA Intern Med. 2019 Jun 1;179(6):820–4.

46 Li D, Kulasegaram K, Hodges BD. Why We Needn't Fear the Machines. Academic Medicine. 2019 May;94(5):623-625.

47 Israni ST, Verghese A. Humanizing Artificial Intelligence. Journal of the American Medical Association; 2019 Jan 1;321(1):29–30.

48 Peterkin AD, Skorzewska A. Health Humanities in Postgraduate Medical Education. Oxford University Press; 2018.

49 Kumagai AK, Richardson L, Khan S, Kuper A. Dialogues on the Threshold. Academic Medicine. 2018 Dec;93(12):1778–83.

50 Kuper A, Veinot P, Leavitt J, Levitt S, Li A, Goguen J, et al. Epistemology, culture, justice and power: non-bioscientific knowledge for medical training. Medical Education. 2016 Nov 15;51(2):158–73.