WITHOUT COMPASSION THERE IS NO HEALTHCARE

EDITED BY: HODGES, BD, PAECH, G. AND BENNETT, J. (2020).

INTRODUCTORY CHAPTER

Technology, Compassion, and the Future of Healthcare

BRIAN D. HODGES

Available for purchase:
McGill Queens University press

THE CHANGING LANDSCAPE OF HEALTH PROFESSIONAL WORK

Emergent technologies will change the future of health professional work , rendering some tasks and professions obsolete.

Healthcare is strongly resistant to change. Despite the dramatic technological changes ahead, healthcare institutions and professions have been slow to prepare. Yet, as a teacher and chief medical officer in one of Canada 's largest healthcare systems , I am regularly asked by students to help them understand what they will be doing in the future. In some cases, they have heard ominous predictions. Professor Geoffrey Hinton of Google and the University of Toronto has famously suggested, for example, that there will be no need for radiologists in ten years.

Simply reassuring the next generation of health professionals that they will continue to have a role is not enough. Students learning to be health professionals today need an education that will prepare them for a rapidly changing context . Much of what their teachers do today will not be what they themselves will do in the future . They must face this truth squarely, as ~~must~~ their teachers – a challenge taken up by Mallette and colleagues in chapter 5 of this book. Current health professionals are also confronting the challenge of practicing as experts while preparing themselves for very different workplaces.

Consider an example. At the Princess Margaret Cancer Centre in Toronto, a new AI-enabled treatment planning system was installed last year, initially for use with breast cancer patients requiring radiation therapy. Formerly, professionals caring for such patients would create a treatment plan to guide the machines that deliver powerful radiation. A treatment plan requires careful calculation of the dose, size, and strength of the radiation beam. Treatment must be located precisely on the patient's body. Incorrect calculations can lead to ineffective treatment or, in rare cases, radiation burns. For decades, a team of three professionals has created those plans: an oncologist, a radiation therapist, and a medical physicist. Together they analyze and integrate data from CT and MRI scans, other tests, and patients' charts. Until recently, the planning process averaged three to four hours per patient.

When I visited the Radiation Medicine Program or Radiation Oncology Department at Princess Margaret Cancer Centre recently, what I saw was amazing: using the new treatment planning system that once took three to four hours now takes about four minutes. Reducing a critical process in our hospital from hours to mere minutes is a tectonic shift in the work of the breast cancer radiation clinic. What happens to the time found through automation? Will that "gift of time" be reinvested into a better patient experience, into high quality and safety, or perhaps simply into the treatment of more patients? Among the hundreds of professionals at Princess Margaret Cancer Centre, there is great anticipation about the promise of emerging technologies such as this one but also a degree of apprehension (Gillan, Harnett, et al. 2018). What will be the future work of these professionals, and how will it change the experience of their patients?

It is sometimes (falsely) presumed that the changes wrought by technology will largely occur in tertiary hospital environments. Yet one of the most compelling examples I have recently encountered is an AI-driven system used in the community at the very front line of care. A simple smart phone or tablet equipped with an app allows community-based personal support workers, chiropodists, nurses, and family doctors to assess a diabetic foot wound. The pixilation of smart phone cameras is now sophisticated enough to analyze the micro-vascular structure (healthiness of the blood supply) of a wound and to determine whether care should be conservative (dressings, antibiotics, etc.) or rather urgent (hospital care) to prevent an amputation. Even in a hospital lab, this analysis for the prediction of healing has never been possible before. Phones and cameras have enabled

the collection and interpretation of numerous data points. This is an amazing advance in terms of bringing care closer to patients. But it will also change which professionals do what and in what sequence to diagnose and treat diabetic foot ulcers.

The rhetoric of technological change can sometimes be overblown. Not all ominous predictions come to pass, and many that do will involve a slower evolution. In the 1980s, when I was a medical student, early electrocardiogram machines provided a rudimentary interpretation of a heart rhythm, but it was not a very reliable interpretation and it didn't displace the interpretation of a physician. Only recently have we developed more powerful computers with machine-learning capabilities that can spot cardiac anomalies just as well as humans can. Whatever the pace, however, professional scopes of practice are changing. The work of health professionals will be reconfigured substantially in coming years. Some functions, and perhaps some professions, may become obsolete even before the current generation retires.

We can predict with confidence that the work of health professionals will change. It is quite another matter to predict *how* it will change. Which roles will be enhanced, transformed, or replaced by machines? As sophisticated as technologies have become in conducting technical tasks, we intuitively believe that some things are so complex that only a human can achieve them. Could a computer ever detect and effectively respond to a patient's fear or sadness? Could it deliver bad news or weigh ethical options? Would we want it to? The next section explores some of these questions.

WHAT WILL WE NEED HUMANS FOR?

The changes to professional work will likely run deep, extending beyond routine tasks. They may disrupt the very foundations of the professions.

Predicting the future is an exercise fraught with risk. Yet science fiction sometimes provides the first accurate sketches. In 1983, a *Toronto Star* journalist asked the futurist writer Isaac Asimov to predict what the world would be like in 2019 (Johnson 2018). Many of his predictions have borne out. He wrote, for example, that computerization would "undoubtedly continue onward inevitably" and that the "mobile computerized object" would "penetrate the home." The increasing complexity of society, he predicted, would make it impossible to live without this technology. Computers would disrupt work

habits and replace old jobs with ones that are radically different while robotics would kill "routine clerical and assembly-line jobs." Further, society would need a "vast change in the nature of education" such that "entire populations" would become computer-literate and able to deal with a high-tech world. Finally, and presciently, he noted that this transition would be rapid and difficult for many.

While Asimov was right in many ways, the pace of our adaptation to technological change in healthcare has been slower than he predicted. He imagined that we would have completed a major transition in our educational systems by 2019, and this has not occurred. He may have underestimated the kinds of jobs that would be affected by emerging technologies as well. Though there is no doubt that much "routine" work has already been replaced by automation, what is becoming apparent is that computerized systems will also impact more complex work such as medical diagnosis, a domain currently reserved for physicians.

In their book *The Future of the Professions* (2015), Oxford professors Richard Susskind and Daniel Susskind caution against the assumption that machines will take over only routine tasks and that human professionals will always be needed for work that is inherently complex. They note that the interpretation of digital images (e.g., X-rays, photographs, scans, or pathology slides) is indeed complex *for humans*. However, such tasks are rather straightforward for computers. A simple *routine* versus *complex* taxonomy is therefore not sufficient to predict what human health professionals will be needed for in the future.

In fact, no one will be unaffected by the coming technological transformations. As Asimov suggested, they may be "difficult for many." This makes it timely to ask: What will happen to the radiologists and radiology technologists, dermatologists and nurses, ophthalmologists and optometrists, oncologists and radiation therapists whose work is displaced, replaced, or transformed by technology?

This book imagines a future in which health professionals are no longer the sole owners of medical knowledge and dispensers of wisdom, nor the most adept or dexterous operators of precision instruments. A profound shift in what knowledge is and how medical procedures are performed is underway. That much is clear. But what about that most human of healthcare domains: compassion? Can we imagine that technology will ever play a meaningful role in listening empathically, understanding deeply, or offering comfort? The next

section challenges a too-tempting dichotomy that humans are compassionate and technologies are not.

HUMANS AND TECHNOLOGIES AS ALLIES IN COMPASSIONATE CARE

In order to understand what kinds of work humans will be needed to do, we first need to acknowledge that humans and technologies are not opponents. Together, they create an urgent need for compassion – and only together are they likely to address that need.

So an AI-enabled computer can read an X-ray. But surely, people say to me, a computer cannot replace human compassion. If the emphasis is on the word *replace*, I ultimately share this conviction. Our case will be built on thin ground, however, if we assume that humans inherently demonstrate compassion and that technologies inherently threaten it. This book is devoted to exploring and championing the importance of compassion from every angle, including how technologies might be able to extend or amplify human compassion. That said, the editors and authors of this book are not uncritical champions of technology: the introduction of deep learning, AI, data analytics, and robots must also include a clear-eyed and critical look at how technology can work against compassionate care, and the book addresses that question too.

Neither are we naïve about the degree to which health professionals and institutions demonstrate compassion today. Though it is safe to assume that most health professionals have the *ability* to demonstrate compassion, professionals and patients will be the first to admit that in busy hospitals, clinics, and communities at large, compassion is often the first thing to evaporate in the push toward greater efficiency. Organizations may commonly cite compassionate care as a guiding value, but – as Martimianakis and colleagues argue in chapter 7 of this book – translating that value into practice requires tangible commitments across all levels of organizational planning. Far from finding consistent compassionate care, encounters with healthcare systems and the humans who run them may actually increase distress or even foster patients' suffering.

For this reason, we must be honest about the need to bolster the compassionate orientation of healthcare professionals and institutions in all forms. This has been an imperative, and a struggle, for hundreds

of years. But as we enter the new technological era, we must begin to ask whether humans could actually collaborate with technologies such as computers and robots to advance compassionate care. While it is essential to understand the uniquely *human* dimensions of compassionate healthcare, it is also important to explore how technologies might be part of the solution.

In advocating for compassion in this book, therefore, we do not situate humans in opposition to technologies. Instead we ask, If compassion is essential in healthcare, *what deployment of human abilities and what technologies together will be most effective*? Each chapter of this book takes up a different dimension of the challenge. Across a variety of domains, the authors ask, What is compassion? How can humans ensure that healthcare remains compassionate in an era of emerging and disruptive technologies?

FLEXIBLE, ADAPTABLE HEALTHCARE PROFESSIONALS

The professions will survive only if they are versatile, focusing not only on knowledge and skills – which will undoubtedly change – but also on the overarching purpose of healthcare. Compassion and human connection are central to that purpose.

New technologies in workplaces of all kinds are challenging our notions about what core elements constitute any job or profession. Industries outside of healthcare provide compelling examples of transformation on a massive scale that resulted from automation more than a decade ago. Successful transformation relies on flexibility and adaptability of the workforce. For example, a recent summary of a report from the Royal Bank of Canada (2018, n.p.) argues that many jobs, even in what were thought to be disparate fields, are in fact connected by a similar set of foundational skills: "Musicians and paramedics might not seem to have a lot in common, but both jobs require high levels of focus, excellent analytical skills, and attention to detail. It takes upgrading only four skills for someone to transition from dental assistant to graphic designer."

While, of course, a great deal of specialized education differentiates a surgeon from a psychiatrist, this line of thinking is valuable to healthcare too. It directs our attention to foundational elements that underpin professional expertise rather than simply specific stores of

knowledge or unique procedural skills. This line of thinking also opens the door to more flexible pathways to competence in the health professions than currently exist.

If a radiologist is no longer needed to read X-rays, is a radiologist no longer needed? Perhaps this is the wrong question. When a radiology resident asks me to speculate about their future, the challenge is to decode the foundational elements of their work that are likely to remain true, even if AI takes over some of the specific tasks that constitute the work today. Rather than thinking of radiologists as specialists who read X-rays, it is more useful to think of them as professionals who use various technologies to visualize the interior of the body for the purpose of medical diagnosis and treatment. This definition helps students and educators to focus on how technologies serve these specific functions and purposes. It also helps radiology residents to focus on learning about the technologies themselves – how they transform the human body into images and, critically, what biases and errors of interpretation they can introduce.

There is another consideration, and this one is crucial. Some of my radiologist colleagues tell me that the future of their profession lies in developing greater focus on the interface with other medical colleagues and with patients. I realized that I had held a stereotype of radiologists as doctors who sit in a dark room reading X-rays, until a colleague said, "I think the future of our profession lies in working closely with medical colleagues and patients to help them determine which technologies to use and how to make sense of the findings." He told me that he had started to work shifts in the emergency department to do just that. "We have to re-emphasize the human interaction part of our specialty."

The twentieth century involved a great deal of work by educators to define what it means to be a competent health professional in different roles. This work took the form of detailed (and often territorial) *competence frameworks* and *scopes of practice* documents. Such documents laid out all the knowledge and skills that students must learn to pass high stakes professional examinations. These frameworks have often been undergirded by Miller's Pyramid, a hierarchical model that articulates competence at four levels: knows, knows how, shows, and does. From this framework an enormous system of assessment and certification has grown, involving professional colleges, accreditation bodies, and certification procedures. It is now almost impossible to pivot from one area of professional expertise to another without

enduring a long program of retraining and recertification. And it is becoming clear that this edifice of health professions education, certification, and accreditation is not suited to the future of healthcare; it does not foster the flexibility and adaptability that the future healthcare workforce will need.

The changing landscape of practice also makes evident that specific domains of sub-specialized knowledge and skills are not enough. Health professionals master a lot of knowledge, including both the memorization of facts and the application of information in practice, such as diagnosis. They also master a bevy of technical skills, ranging from the simple taking of blood pressure to the most complex of surgical procedures. Yet for all that knowledge and skill, the practice of health professionals, regardless of specialty or profession, is held together by a remarkably similar core purpose: to care for other humans in ways that alleviate suffering. That is healthcare's compassionate purpose. It is therefore alarming to note that some competence frameworks have lost sight of compassion altogether. Research supported by AMS Healthcare in the last decade has shown an almost complete disappearance of the word *compassion* (and, more astoundingly still, the word *care*) from the competence frameworks in medicine and nursing (Whitehead, Kuper, et al. 2014).

While specialized knowledge and skill will always have great importance, as computers develop abilities that rival human brains, the interpersonal domain may become our greatest differentiator. From a strictly economic perspective, "competencies that are complementary to machine prediction will become more valuable in the future, while competencies that are substitutes for machine prediction will become less valuable" (Li, Kulasegaram, and Hodges 2019, 623). In many ways this is a positive realization; it recognizes that our human capacity for compassion is much less specialty specific than is our cognitive knowledge – and that without it healthcare cannot exist. Put another way, all the knowledge held in internet clouds combined with the most sophisticated pattern recognition of AI, bolstered by the dexterity of the latest robots will not be sufficient to create high-quality healthcare. These technological capabilities, extraordinary as they are, must be combined with the human capacity for compassion in order to deliver the complete healthcare package.

But in order for healthcare professionals, educators, and patients to fully understand how human compassion will fit into healthcare systems of the future, we first need to take a closer look at what is

happening to the domains of professional knowledge and skills that advanced technologies are disrupting. We must consider what health professionals may no longer be doing.

KNOWLEDGE IS IN THE CLOUDS

Well-established factual knowledge, once central to professional expertise, is readily replaced by computers.

The television show *Star Trek: Voyager* gave us the first glimpse of a holographic doctor. In creating the character of "The Doctor," the writers anticipated what an AI-enabled computer diagnostic system might be like. Though a real actor played The Doctor, it strikes no one as odd today that the knowledge necessary for the practice of medicine could be stored in a non-human database activated by voice. In fact, most of us already use this technology every day in our phones. You can ask a smartphone, "What is a normal level of blood potassium?" or "Does penicillin interact with grapefruit juice?" and expect to get a meaningful answer. For this reason, there is little justification for health professional students to memorize thousands of pages of biochemical formulae, anatomical parts, or drug names and interactions. Memorization has become, in the view of many, a waste of cognitive resources. Most of medicine's factual knowledge is not held in human brains; it is accessed from computer databases.

Databases of factual knowledge also enable patients to be more participatory in their own care. Many people, with and without professional training, consult the internet to learn about medical problems. As a physician, I regularly use the internet to find evidence related to the treatment of my patients, but I also use it in relation to my own health. Recently, I found a video that helped me to treat my sprained ankle. (It has been a long time since I completed my generalist training!) Patients and families can access the same information as health professionals. Imagine that you develop repetitive strain injury. Are you likely to go through the trouble of getting an appointment with a healthcare professional if you can easily access and understand information provided online? I recently had just that experience and found an excellent video of exercises to counteract the repetitive pain I get from too much typing.

Such well-established factual and procedural knowledge – similar to what the ancient philosopher Aristotle termed *episteme* – has

traditionally informed the central dimensions of professional expertise. In the foreseeable future, a role will persist for specialized knowledge and for the interpretation of complex signs and symptoms. This role, however, will continue to diminish. Increasingly, databases can be accessed for the purposes of analysis and learning directly by computers, such that computers can master more extensive and complex forms of knowledge. They are also able to provide informative probabilistic analyses of diagnostic signs and symptoms. Machine learning and AI will continue to evolve rapidly and will be widely used to support human interpretations and judgments of patterns of illness and disease. Sometimes, people will bypass health professionals and go directly to the internet for interpretation of symptoms or recommendations for treatment.

This suggests an emerging role for human health professionals. Much information on the internet is of poor quality, and there will be a role for professionals to help patients deal with the clouds of facts, many of which are replete with biases. Even those facts that are effectively uncontested need to be appropriately selected and applied. Human professionals will be needed to help people navigate, evaluate, and interpret competing perspectives and available options. Eric Topol, author of a landmark report to England's National Healthcare Service, observes, "The new medicine envisioned will require extensive education and training of the clinician workforce and the public, with cultivation of a cross-disciplinary approach that includes data scientists, computer scientists, engineers, [and] bio-informaticians, in addition to the traditional mix of pharmacists, nurses and doctors" (Topol 2019b, 6).

Perhaps the first task for this large and diversified healthcare team is to understand much more about the differences between humans and computers in how they think and what kinds of errors they're prone to make. Only then will we know how humans and machines can work most effectively together.

RECOGNIZING PATTERNS: SKIN LESIONS, CHIHUAHUAS, AND BLUEBERRY MUFFINS

Computers and human experts each excel at pattern recognition in different ways – and each is vulnerable to different kinds of errors. Humans are ultimately responsible for understanding the specific power and limits of technologies, and for deploying them appropriately.

Humans have highly developed abilities to take in thousands of pieces of diverse sensory information and rapidly form an impression. Health professionals learn to do this quickly from seeing hundreds of patients, and they test their initial impressions by asking careful confirmatory questions and analyzing physical signs and laboratory results. Consequently, health professionals become experts at recognizing patterns.

Computers are also very good at pattern recognition. Computer-based pattern recognition involves digital algorithms programmed by humans. Increasingly, in the case of machine learning, computers themselves will also be able to generate and modify algorithms as they interact with large data sets. Hundreds of such algorithms already operate behind the scenes in our daily lives.

Algorithms are nothing new. They have existed as long as there have been computers and were used by humans long before that. A decision tree that helps a doctor assemble signs and symptoms into a diagnosis is a simple algorithm. What has changed is the sophistication of algorithms and the speed at which they can be automated and employed. Whereas algorithms were once of interest only to computer programmers, today they govern our daily lives. When we interact with the internet, algorithms make analytic judgments *about us*, by comparing our personal qualities or features against a database. For example, search engines such as Google use algorithms to determine what advertisements to make visible to us. An algorithm incorporates data about everything we have done and shared online: past searches, purchases, personal demographic data (often harvested from social media), where we live and shop, our age, gender, and culture. It then makes predictions about what we will be interested in, what we are likely to click on, and what we might purchase. The algorithms are used to push customized information to us.

This is also how algorithms work in healthcare. Many pieces of information about you – including your medical history, CT and MRI scans, X-rays, blood samples, biopsies, and psychological tests – can be compared to a huge database of other people's information. By comparing your data to norms, computers predict the presence or absence of diseases and probabilities of future events

Yet pattern recognition is neither completely objective nor neutral. Both humans and machines can make errors, though they tend to do so in different ways. Take, for example, a popular meme that illustrates the challenge for AI of distinguishing a chihuahua's face from a

blueberry muffin. If the two little eyes and nose are similar in configuration to three blueberries on a muffin top, a computer can confuse the images. This example is perhaps something of an urban legend now; AI can generally (though imperfectly) pass this test. The nature of this error nevertheless illustrates an important point: a five-year-old human child would not confuse a dog and a muffin.

An interesting question, as healthcare moves forward with machine-human collaboration, is whether computer-generated predictions will be able to support patients in ways that are qualitatively similar to the guidance of human health professionals, or whether they will simply mimic that support in a superficial fashion. Will Dr Google's advice be simply watered-down medical care or could it contribute a valuable new dimension? One area of potential added value is prognostication. Humans can easily miss subtle patterns in layers of complex patient data. A study aiming to predict survival in patients with heart failure, for example, showed that an AI-enabled computer was better able than physicians to integrate the data from scans and tests with eight years of chart data, resulting in more accurate predictions of survival by the computer than by the physicians. This is perhaps unsurprising, given the cognitive difficulty for a human of amassing and integrating so much information. No physician has time to thoroughly read eight years of charts.

So humans and computers each have strengths, weaknesses, and blinds spots in performing pattern recognition. What is critical is that we understand (and teach) the specific biases and types of errors to which human and computer processes are prone. These biases and errors may be radically different. To avoid them, health professionals in the future will need to have a much better understanding of how human and artificial minds work, and how they work together. It will not be sufficient in healthcare to use information technologies in the way that we tend to use search engines like Google – entering questions into a blank box, which uses a process that we don't understand, and simply accepting the answer as correct.

But it is also fallacy to believe that human judgment is free of bias. Many studies confirm that humans are prone to all sorts of distortions in recognizing patterns and making judgments. Some derive from transient human weaknesses like fatigue or distraction. Others derive from ingrained cognitive biases, such as the halo effect of interpreting new information through old assumptions, or the recency effect of being influenced by the case seen just before a new one. Still others

derive from pervasive social biases, such as the stereotyping of racial, linguistic, gender, and cultural factors. Human judgment is replete with bias. What is becoming visible is a new concern that when humans build computers, we may create algorithms that actually amplify human biases.

In her book *Algorithms of Oppression*, Safiya Umoja Noble (2018) illustrates this problem with a simple example. She shows the very different information that was returned from searches of the internet with the terms "Black girls" versus "white girls." The former search returned a high number of pornographic and racist sites, while the latter produced such things as preppy college websites and beauty products. These differences could be explained because the algorithms that determined associated terms were derived from other past searches. In other words, the search algorithms built on and amplified very human, and in this case racist, biases. Similar research in the social sciences reveals that algorithms can embed discrimination when they are used to determine who is eligible for social programs such as housing or employment benefits (Eubanks 2018). In chapter 3 of this book, Paton and colleagues build on examples such as these to show how compassion is interconnected with the broader concept of equity.

This points to an important role for human professionals who wish to use new technologies that provide diagnoses and prognoses in a way that is compassionate. Far from assuming that computer-based systems will be more objective, human health professionals will likely have to be even more vigilant in ensuring that diagnostic systems are fair, accurate, and objective.

WISDOM GOES BEYOND PATTERN RECOGNITION

Computerized algorithms can be insensitive to cultural and situational specificities. Humans should strive not only for practical knowledge, which involves recognizing and applying patterns but also for wisdom: judgment in specific situations that integrates factual and technical knowledge with ethical and interpersonal sensitivity.

When a doctor or nurse interacts with an algorithm-driven system to determine if a patient's test results are normal or to predict their clinical outcome, it will matter if they also know what data were used to create the database, what the algorithms are looking for, and what sorts of erroneous assumptions or errors might be made.

Such issues are pressing today because healthcare applications using AI are already coming online. The Food and Drug Administration in the United States has recently licensed for clinical use an AI-driven diagnostic system that can classify skin lesions (FDA 2019). It is easy to take a picture of a new brown spot on your arm and ask the system the likelihood that it is a freckle, a benign mole, or a cancerous melanoma. The system is accurate and will be a boon to both patients and health professionals, particularly given how hard it is to access a dermatologist. To perform this remarkable task, the system had to be trained. Training involved teaching the AI system to recognize skin lesions by showing it pictures and telling it the right answer. The system became quite reliable after it had seen nearly 130,000 images together with the correct diagnosis (Esteva, Kuprel, et al. 2017). The AI system learned to rapidly and accurately identify skin lesions it has never seen before and to triage patients into high-risk or low-risk groups for follow-up. But what is the risk of such applications? The pictures used to train such systems primarily represent Caucasian patients. How will such systems perform when classifying skin lesions from people with other skin tones (Lashbrook 2018)?

At a higher level of complexity than skin lesions, AI systems are available to recognize human emotions and behaviours. Some of the inventors would have us believe that computers even have the power to "eliminate human bias" in interpreting patterns of human behaviour. Brown (2107) describes how emotion recognition technologies detect "subtle 'micro-expressions' associated with joy, trust, fear, surprise, sadness, disgust, and anger" in order to quickly and accurately predict people's emotions and motivations. Such descriptions beg careful consideration of how those characteristics were defined and compared to what set of "normal" data. These limits and biases often remain hidden from view. We surely cannot call this technology accurate or indeed compassionate if we believe, as I do, that compassion requires authentic understanding.

A skin lesion, serum potassium measurement, or picture of a retina presents specific questions and concrete data for computer analysis. By contrast, most visits to a health professional are initiated by more ambiguous patient complaints or problems. Unlike reading an image, understanding a complaint (or set of symptoms) is more involved than simple visual pattern recognition because it requires attention to physical, physiological, biochemical, social, and psychological elements all at once. This integrated understanding is a key element in

compassionate care. If a clinician is to generate understanding with a goal of relieving suffering then they must employ careful observation and skilled inquiry to interpret multi-dimensional data. That data may also include what is not said at all. Consider a woman who complains of abdominal pain and has subtle abrasions on her arm – symptoms that make little sense until the physician or nurse recognizes that she averts her eyes, a clue to ask about domestic violence.

Even health professionals who are skilled in the use of newer technologies voice caution about their application in healthcare. Dr Isaac Kohane, Chair of Harvard Medical School's Department of Biomedical Informatics, is skilled in the application of "big data." Yet in an article titled "The Beauty of 'Small Data' in Medicine" (2018), he shares a cautionary tale from his training. Kohane recounts how he met a nine-year-old boy referred to his clinic for short stature. His first impression was that the child didn't seem very short, though a family photo showed that the child was shorter than his seven-year-old brother. Normal practice would have been to plot the boy's growth, perhaps order an X-ray and blood tests, and send his family home with reassurance. But Kohane's teacher, a man with great experience in child development, noticed a subtle abnormality on the growth chart of the child that prompted him to repeat the measurements and recommend a brain scan. The scan revealed a benign brain tumour that was (happily) removed by a neurosurgical team. Dr Kohane's message? The case "marked the beginning of a long education on the value of small data – that is, the clinical impact of a small number of reliable measurements on a single patient." He notes that his own role as a "big data practitioner" offers all the more reason "to remember how much can be done with careful, meticulous consideration of data coming from a single patient" (Kohane 2018, n.p.)

Compassion requires such close attention to the symptoms and experiences of each individual. It goes well beyond the simple application of an established database to recognize a pattern. Integrating signs and symptoms with clinical experience, practical wisdom, and an ethical framework is far more complex. Compassion is not only specific to individuals but also to cultures. Eliminating biases might be highly desirable from the point of view of reducing superficial diagnostic assumptions and errors; however, what does bias mean in sophisticated, culturally determined human behaviours and emotions? Consider eye contact. Many standardized checklists of communication skills start with the item "makes eye contact" at the top. Such

checklists constitute simple algorithms for judging trustworthiness. Humans attach significant meaning to the act of making or avoiding eye contact. If someone averts their eyes, others often suspect that they are hiding something. This interpretation is built into algorithms used at airports to detect suspicious travellers. How long until a computer-assisted diagnostic system includes eye contact within assessments of depression, anxiety, or perhaps even truthfulness concerning personal relations or drug use?

Yet, in Indigenous cultures, making eye contact is considered rude. I have worked with Indigenous peoples in Canada's Arctic communities who tell me that direct eye contact with a stranger feels like inappropriate touching. For some Indigenous peoples, and indeed in many cultures, avoiding eye contact is a sign of respect. If this interpretation of eye averting is not captured in an automated system (perhaps because few Indigenous people were part of the database), the system would have a built-in bias that could lead to significant problems of interpretation and of understanding.

Because compassion is in part about understanding, any technology that we imagine to have compassionate uses must be evaluated through a socio-cultural lens. When health professionals use machines to aid in diagnosing a skin lesion or assessing mental capacities and behaviours, they must be vigilant about what algorithms include, what they leave out, and to what databases and norms they are compared. Further, as algorithms begin to help control who can and cannot access healthcare and social services, very human biases related to gender, race, religion, and other socio-demographics may too easily be built in. Ultimately, it must be humans who maintain vigilance to ensure ethical and compassionate uses of technology.

WHEN COMPUTERS LEARN THE MEDICAL GAZE[1]

Technologies have introduced distance between healthcare providers and patients. This distance has often been a threat to compassion. Health professionals need to be aware of this distance and find new ways to foster presence.

All medical students are shown the picture in figure 0.1. It is in some ways the primal image of medical care. A caring doctor gazes at a poor, sick child, while a mother weeps and a stoic father looks on with concern. There are some stereotypes at play here, but in essence,

Figure 0.1 *The Doctor* by Luke Fildes (1891) illustrates the medical gaze

this is how many continue to imagine that a healthcare professional brings compassion to patients. Philosopher Michel Foucault, in his book *The Birth of the Clinic*, describes this *gaze* as a knowing and penetrating way of looking at a patient. *The Birth of the Clinic* traces the origin of the concept of clinical medicine itself. Though written a half century ago, it continues to serve as a relevant history of the ideas that underpin healthcare. Prior to the eighteenth century, the practice of medicine was largely mystical: theories abounded of misplaced organs, blocked humours, wind in the joints, and flows of energy. The rise of scientific medicine, with its dissections, experiments, and animal studies changed all that.

But Foucault notes another change in this seminal work. He cautions that this gaze – the objectifying, apparently neutral, scientific, medical gaze – could, in fact, turn patients into objects of study, rather like the way one looks down a microscope at an insect. Today, his work helps to explain how health professionals sometimes drift from

caring for patients to studying them as objects. Sometimes, in our best attempts to help people, we stop seeing the people and begin to focus more on the diseases, the cells, or the X-rays that come from them.

The danger of objectification began long before the rise of AI and robots, but the objectifying medical gaze becomes more powerful still when a machine does the gazing. While a human health professional may be prone to inattentiveness, we can return to awareness of our patients' humanity. There is no such awareness for an automated system. Human health professionals must ensure that patients' experiences are recognized and acknowledged – that their treatment is *humane*, as they interact with increasingly non-human technologies.

Some technologies will develop superficial forms of empathy. Some might even seem polite or convincing. Indeed, all machines should be designed with caring as one goal of the experience for people who use them. In chapter 1 of this book, Wiljer and colleagues begin the important work of mapping how technologies can mediate compassionate care in a variety of ways across healthcare ecosystems. Nevertheless, human presence will never be entirely replaceable: a machine alone can never have a compassionate aim. One main function of human healthcare professionals in the future may be to recognize the power that simply being present has to reduce the objectification of patients. Judith John is a long-time patient activist at Toronto's University Health Network. She advises new health professional students and seasoned veterans alike: "I want you to be present for me" (John 2016).

The etymological root of the word *compassion* is "suffering with," which suggests that compassion requires human *presence* – a principle central to the discussion of patient engagement in chapter 2 of this book (by Rowland and Johannesen) and of compassionate leadership in chapter 6 (by Tassone and colleagues). If we value presence, it becomes problematic when technologies distract, displace, or diminish humans in their connections with each other. A simple example is the rise of people walking around the streets with their eyes fixed on their mobile phones. This common behaviour, reinforced by a perceived need for (or addiction to) *virtual* interaction, clearly interrupts human interaction in the physical world. Similarly, when a clinician turns her back on a patient to enter data into a computer, the human connection is partly broken. Leaving aside for the moment the possibility of

technologies that can reinforce human interaction, it seems self-evident that patients will want to be known, cared for, and comforted by other humans – to feel their presence and to benefit from their attentiveness, at least sometimes in the course of their care. It also seems self-evident that many, if not most, who work in healthcare are governed by the value they place on human interaction. Indeed, my personal suspicion is that the current epidemic of burnout among health professionals is largely driven by the diminution of human contact in our work environments. In chapter 4 of this book, Maunder and colleagues delve into supporting evidence: relationships are integral to fostering resilience and protecting against burnout. Healthcare work will face a mounting crisis if there is continued erosion of human presence and human interaction.

THE DISTANCE BETWEEN US

Amazing new technologies enhance clinical care, but they are often interposed in the physical space between health professional and patient. This can interfere with the human connection.

The *laying on of hands* has traditionally strengthened the relationship between patients and health professionals by creating a physical bond. There was a time when the physical examination was the core of medical diagnosis. When I trained in medical school thirty years ago, we learned about things such as "whispering pectoriloquy" (the sound of the whispered voice heard through a stethoscope) and diaphragmatic excursion (the movement of the diaphragm measured by tapping, or "percussing," with one's finger along the chest wall). Today, my students say, "Why would we do that? Don't we just get an X-ray or a 2-D echo?" Something has been lost, I think, in adopting these admittedly more accurate diagnostic tools. The physical exam was not about the diagnostic process alone. The physical contact also established human connection.

About two centuries ago, a physician in France named René T.H. Laënnec invented the stethoscope. Indeed, the stethoscope is a marvellous tool. From the first one – which was just a rolled up paper tube placed against the patient's chest – our diagnostic capability has been dramatically improved. With a stethoscope, a physician, nurse, or respiratory therapist can hear breath sounds, heart sounds, and bowel

sounds with ease. A little later, a modification allowed physicians and midwives to hear the rapidly beating heart of a developing baby in a woman's abdomen.

The stethoscope, though, created a little physical space between the physician, nurse, or midwife and the people they examine. That small space has been growing wider and wider. Today the stethoscope is largely a historical object. Though many health professionals still carry one, it functions more as a symbol of a professional role rather than a tool in actual use (Bernstein 2016). Of course it still functions, just as it always did, but most of us now are thinking more about the human relationship when we use one.

I am mindful of when and how I touch my patients. As a medical psychiatrist in a large hospital, I work in intensive care units, on organ transplantation services, and in the emergency department. The patients I see are often confused or delirious. Some come to hospital with a psychiatric disorder such as depression or anxiety. It is really very difficult for anyone to be in a place like an intensive care unit. The lights are never turned off; the machines, such as respirators and intravenous pumps, whir and beep constantly. The healthcare professionals circulate on rounds at all hours of day and night. If I stand at the bedside in a starched white coat wearing a mask and only use my voice to communicate with a patient attached to all those machines, we are both cut off from the human relationship we need to have. So, I always place my hand on a patient's shoulder, hold their hand, or cradle their wrist, perhaps going through the motions of checking a pulse to normalize the touching. Thus a human connection is established.

When my own appendix ruptured in 2010, I was amazed to experience as a patient how little human contact there was when the diagnosis was made and communicated. My doctor reported to me, "The CT scan showed that you have a perforated appendix." I thought, Wow, the *CT scanner* made my diagnosis? Of course it was a human radiologist who examined the CT image, made a diagnostic conclusion, and called the emergency physician. But I never met the radiologist, and the emergency physician performed only a cursory examination of my abdomen. I had no actual connection with either.

Technologies can help with diagnosis and treatment, but they can also create new distances and barriers to human presence. This is particularly true of the most technologically intensive areas of care, such as surgery. Health professionals of the future working in these

environments will need to find strategies, old and new, to forge human connections across increasing distances.

AUTOMATION AND TECHNICAL SKILLS

Technologies greatly extend the technical precision of humans. In surgery, these technical capacities are still deployed by human judgment – and patients tend to place their trust in humans.

While the cognitive domains of healthcare, including pattern recognition, diagnosis, and prognosis are undergoing transformation, there is an equally profound shift underway in technical skills. Consider how technology is changing brain surgery. At Toronto Western Hospital patients now commonly arrive in the morning, have a craniotomy (the skull opened or small holes created), then wires or shunts implanted or a tumour removed, all in time for them to return home the same evening.

How is it possible that performing a brain implant or removing a tumour has become same-day surgery? Many technologies have contributed to transforming what were once long, dangerous operations with many days in hospital into one-day procedures. Among them are improved surgical instruments, better sterilization, high-quality imaging, precision lenses for neurosurgeons, and anesthesia that removes all feeling of pain but allows patients to remain awake. Indeed, the success of all forms of surgery has dramatically improved in the last century. Today, aided by tiny rods and cameras, surgeons can operate in minute spaces where clumsy hands would do damage. These technologies are arguably compassionate because they improve outcomes while vastly decreasing the suffering of patients.

Even more dramatic is the arrival of surgical robots; they are extraordinary to watch. Though some robots look humanoid, with arms, legs, and heads, surgical robots are not like that. In fact, most robots do not look like humans. While mimicking human anatomy may have some advantages, we humans are actually rather limited in a number of ways. Human hands, while amazingly dexterous, for their size they do not work efficiently in small body cavities. Enter the surgical robot.

The first time I watched a surgical robot in action was in a gynecological operating room at Toronto General Hospital. While I had spent many weeks in surgery as a medical student and during my

generalist training, it had been years since I had visited an operating room. I was immediately struck by the transformation. A woman (the patient) was suspended, anesthetized, from the ceiling by a series of straps and harnesses. One doctor (the anesthesiologist) sat near her overseeing the tubes delivering the sedation. Another doctor (the surgical assistant) stood by, observing as the metal rods of the robot moved in and out of tiny holes in the patient's abdomen. The surgeon herself was seated at a console across the room, facing a screen on which she could see inside the patient while she operated hand and foot controls. This is not the operating room of my training or that most people imagine.

Humans collaborating with machines such as surgical robots are changing what it means to work as a health professional. Robots are fast becoming our *team members*. Healthcare will require humans to do things, but increasingly machines are augmenting our steadiness, precision, reach, and accuracy. In the examples I have described, robotic systems take up some of the direct, technical work while humans use their judgment to deploy and monitor those systems. The impact on outcomes, such as healing and recovery time, is tremendous and positive. However, in the process, a new risk presents itself: the physical distance grows between professionals and the people they treat as technologies become intermediaries between clinician and patient.

Author of *The Digital Doctor* Robert Wachter remarks that healthcare's path to computerization "has been strewn with land mines, large and small. Medicine, our most intimately human profession, is being dehumanized by the entry of the computer into the exam room." He adds, "While someday the computerization of medicine will surely be that long-awaited 'disruptive innovation,' today it's often just plain disruptive: of the doctor-patient relationship, of clinicians' professional interactions and workflow, and of the way we measure and try to improve things" (Wachter 2015, xi).

In addition to the effects of growing physical distance, there is another consideration when machines join the healthcare team: they can fail. Who hasn't had the experience of a program crashing, or a computer rebooting to update software, at a critical moment? Such failures can cause human emotions to flare in machine-human encounters. But the stakes are even higher when machine failure arises in healthcare settings that are already characterized by emotional tension. What happens when a surgical robot reboots or a radiation therapy

machine responds unpredictably or fails outright? A burgeoning literature documents the rising problem of failing medical devices (Ferrarese, Pozzi, et al. 2016; Hengstler, Enkel, and Duelli 2016). The radiation treatment team at the Princess Margaret is very aware of an infamous and devastating radiation machine failure that led to several deaths in the 1980s. As Jamie Lynch (2017, n.p.) remarks in a blog post titled "The Worst Computer Bugs in History," while such cases are extreme and rare, "they are worth studying for the insights they can offer into software development and deployment. These computer bugs left a significant impact on the people who experienced them, and we hope they'll offer valuable lessons we can all apply to our own work and projects."

A key issue in the evolution of human-machine technical skills, then, is learning to grapple with very human responses – the "significant impact on the people who experience them." An emerging phenomenon is called *computer rage* or *tech rage*, "an overwhelming emotion caused by frustration with one or more technological devices" (Shaw 2015). Neither computers nor robots are autonomously motivated by a sense of urgency. Nor do they feel fear or panic. Thus it falls to the humans who acquire, maintain, and operate machines to recognize and modulate their own emotions *and* those of the patients they are caring for when problems with technology arise.

Unfortunately, the challenge of dealing with the stressors of human-machine interaction and the problem of interpersonal distance between health professionals can converge. Consider a homecare nurse who cares for patients in the community and uses a telehealth system, a radiologist who works in her office in one country and interprets a CT scan taken in another country, or a gastroenterologist in a hospital who uses a machine-guided system to perform a biopsy. All of these professionals benefit from significant gains in automation and autonomy and from distributed models of care. But all of them work in ways that are more isolated and further removed from the colleagues with whom they can confer. A surgical fellow, shaken, recently told me that he experienced three different pieces of equipment failing during a critical operation. Though surgical fellows are fully qualified physicians, they are learning a subspecialty and still require support and supervision from experienced experts. In this case, a number of operating rooms were running simultaneously and the fellow was unable to reach his supervisor. Marshalling his creativity, with input from the nurses and anaesthetists, he managed to stop a serious

hemorrhage. But he told me afterwards that it was all he could do to control his rage at the malfunctioning equipment. This fellow recognized that his emotions were significantly impinging on his problem-solving in the operation room. As he told me darkly, "Of course I can get someone by email or text, but that doesn't help any of us when all hell is breaking loose."

VIRTUAL EMPATHY AND COMPASSION?

Humans have a strong need for close emotional and physical contact. Mediated and virtual connectivity may impoverish human connections, becoming harmful when clinicians and patients need emotional support.

Research among our primate cousins shows that isolated individuals will die from a starvation for contact, and loneliness among humans is associated with increased mortality (Holt-Lunstad, Smith, et al. 2015). Many of today's new technologies, including the varied forms of social media, bring the promise of greater connectivity. Indeed, communication channels such as email, text, or social media can bridge distances and help form new connections. In healthcare, the rise of the "virtual" visit conducted over email or videoconference can augment accessibility to care. But we need to be careful with these technologies and not assume that they automatically confer the same value for human contact.

I saw a young man in my office recently for symptoms of depression. When I asked about his relationships, he told me that he had 107 friends. Of course, those were friends on social media. When I asked him how many people in his life he could trust, talk with face to face, and discuss his feelings with, the answer was zero. I learned that he had great difficultly talking casually with people at school or at the gym, and he found he was only able to communicate through the mediation of his mobile phone. I spent weeks helping him, little by little, learn the basics of sharing casual conversation with strangers and eventually meeting some new people "live." This is not to say that everyone suffers from interpersonal challenges as serious as this young man's, but there is a general trend to relate to one another via communication technologies, and the consequences of that trend warrant careful consideration.

As a physician, I provide support to my patients by email or text. Many of them find this comforting. However, I am aware of some early research revealing the limits of mediated communication in my

own field. When our Department of Psychiatry at the University of Toronto began video consultations to the northern parts of Canada three decades ago, it quickly became clear that appointments via video conference were much more effective if the clinician and patient had already formed a personal, face-to-face relationship.

In some cases, patients may prefer virtual alternatives to traditional forms of communication. Talking with a human to book a clinic appointment, for instance, may be less desirable than having a good app that allows easy appointment reservations and automatic confirmation. Using a robust online system to access preoperative information may be preferable to driving across town, or farther, to hear it from a person at the hospital. On the one hand, many people already benefit from and appreciate the ability to see their laboratory results online in patient portals. On the other hand, non-human systems for such things as counselling, psychotherapy, or the delivery of bad news – all of which are being developed – may garner mixed feelings and results for patients. Among health professionals themselves, email and text are not the most supportive mediums for communicating in a crisis situation.

As we parse the value of different technologies in their ability to augment or detract from human presence, a nutritional analogy comes to mind. Diet soda looks and tastes much like food, but it's an illusion. There is no nutritional value in diet soda. It seems to me that many forms of communication in healthcare today provide a diet soda version of compassion: they appear to foster human relationships, but the value of the resulting connection is more illusory than real. Not all live human interactions are rich in compassion, of course. But the risk grows for meaningless communication and *pseudo-empathy* as the medium becomes more depersonalized. Most of us have experienced the phenomenon of receiving (or sending) an email message that we would never consider appropriate in face-to-face communication. Technological mediation of communication, and ultimately its connection to empathy and compassion, is complex. Mediated communication is not necessarily bad, but it does require good design. In healthcare, it is essential to understand how to build and sustain relationships that foster empathy, understanding, and ultimately compassion, regardless of the medium.

To function effectively, to support each other, to stay calm when there are problems, and to provide care that is compassionate, healthcare professionals need to be savvier at using mediated forms of communication with patients and with each other. Attention will have

to be paid to limiting formats that compromise understanding or diminish empathy. Importantly, health professionals must pay attention to the effects of different technologies on their own well-being.

Dr Atul Gawande is a gifted medical writer and self-confessed technophile. Yet he draws strong links between the uptake of computerized processes and the burnout of health professionals: "Something's gone terribly wrong." He remarks, "Doctors are among the most technology-avid people in society; computerization has simplified tasks in many industries. Yet somehow we've reached a point where people in the medical profession actively, viscerally, volubly hate their computers" (Gawande 2018b, n.p.). While he holds a clear-eyed view that the tools available today greatly enhance our ability to collect, store, and analyze information, leading to better and safer diagnoses and treatments, he notes that much can be lost along the way. "The story of modern medicine is the story of our human struggle with complexity. Technology will, without question, continually increase our ability to make diagnoses, to peer more deeply inside the body and the brain, to offer more treatments. It will help us document it all – but not necessarily to make sense of it all. Technology inevitably produces more noise and new uncertainties" (Gawande 2018b, n.p.).

REDOUBLING THE COMMITMENT TO COMPASSION IN HEALTHCARE

Compassion has always been an anchor of healthcare. The technological revolution before us is a burning platform to restate and reinvigorate the commitment to compassionate healthcare.

As we have seen, all the knowledge and skill in the world are not sufficient to qualify as a good doctor or nurse. One needs years of experience to know when and how to deploy knowledge and skills in ways that enable compassionate care. This is true whether the knowledge is held in a computer database or the skills involve a robot. There are many elements of practical wisdom including judgment, reflection, and adaptation to context. And there is no quality more central to good healthcare than compassion. Myriad studies demonstrate that human relationships affect health outcomes: communication, empathy, and ultimately compassionate healthcare are all related to better treatment adherence, reduced anxiety, increased trust, less need for pain medication, and even better rates of recovery and survival.

It is time to take a hard look at how humans can work together with technologies to enhance these outcomes even further. But as the many examples I have discussed illustrate, this "working together" is not as straightforward as acquiring and installing a new computer or robot.

Consider one last example. Perhaps you saw the movie *Her*? Actress Scarlett Johansson plays the role of the computer operating system that speaks to the main character, Theodore, played by Joaquin Phoenix. When the film opens, Theodore has had a rough time: he is isolated, unhappy in his job, and has lost his partner. The computer operating system asks, "How are you feeling today?" and suggests, "I notice you didn't go to the gym this morning. Maybe you'd like me to arrange for you to go for dinner." We're not so very far from this today. The ability of AI to ask probing questions and to respond is rapidly advancing.

Do you think of this as compassion? Did the creator and operator of the technology in this fictionalized situation have a compassionate purpose? On one hand, Microsoft creator Bill Gates said, "Technology is unlocking the innate compassion we have for our fellow human beings" (Gates 2013). An interesting thought. On the other hand, the Dalai Lama said, "I think technology really increased human ability. But technology cannot produce compassion" (Almendrala 2014). The editors and authors of this book see in these remarks a dated dichotomy. It is not a simple competition of human versus machine. We believe that a compassionate healthcare system is one in which gifted humans and the best technologies collaborate to create exceptional, compassionate care.

As technology advances relentlessly, even rapaciously, healthcare is transforming radically. This book is therefore a call to action. We call upon healthcare professionals, leaders, educators, policy-makers, patients, and families to act: to shape a future in which healthcare is effective, accessible, efficient, and also fully anchored in compassion.

NOTE

1 I use the word *medical* in the broadest sense to indicate the field of practice of medicine and all the health professionals engaged in it. For the specific role of medical doctor I use the term *physician* to avoid confusion, as health professionals of many kinds have a doctoral degree (PhD or DPhil) and may use the title *Doctor*.