Medicine, Health Care and Philosophy (2025) 28:339–349
https://doi.org/10.1007/s11019-025-10264-9

SCIENTIFIC CONTRIBUTION

The need for epistemic humility in AI-assisted pain assessment

Rachel A. Katz[1,2] · S. Scott Graham[3] · Daniel Z. Buchman[2,4,5,6]

Accepted: 27 February 2025 / Published online: 15 March 2025

Abstract
It has been difficult historically for physicians, patients, and philosophers alike to quantify pain given that pain is commonly understood as an individual and subjective experience. The process of measuring and diagnosing pain is often a fraught and complicated process. New developments in diagnostic technologies assisted by artificial intelligence promise more accurate and efficient diagnosis for patients, but these tools are known to reproduce and further entrench existing issues within the healthcare system, such as poor patient treatment and the replication of systemic biases. In this paper we present the argument that there are several ethical-epistemic issues with the potential implementation of these technologies in pain management settings. We draw on literature about self-trust and epistemic and testimonial injustice to make these claims. We conclude with a proposal that the adoption of epistemic humility on the part of both AI tool developers and clinicians can contribute to a climate of trust in and beyond the pain management context and lead to a more just approach to the implementation of AI in pain diagnosis and management.

Keywords Artificial intelligence · Pain · Bioethics · Predictive analytics · Epistemic humility

Introduction

Pain, specifically pain measurement, has historically been a fraught topic in both philosophical and medical circles. Given the many challenges of pain measurement, the process of diagnosing, and treating or managing pain is frequently long and arduous, and that process is made more difficult in the context of longer lasting or chronic pain. The issues surrounding pain are multiple and complex in nature. While the development of new artificial intelligence-assisted medical technologies seeks to simplify these issues and improve diagnostic specificity and accuracy, there are multiple areas where this technology can potentially lead to ethical issues in understanding and treating pain.

Pain measurement is rife with combined epistemic and ethical complexities. Many of these challenges arise from the lack of reliable pain biomarkers, a situation that leads to physicians routinely underestimating patients pain (Seers et al. 2018). This regular conflict between physician assessment and patient report creates an experience where many patients feel as though their physicians do not believe their testimony of being in pain, or if they do, do not take it seriously. Importantly, the gap between patient report and physician assessment is greater for marginalized and minoritized patients, and this has led to significant inequities in both clinical care and patient outcomes (Mathur et al. 2022). Some of the current ethical-epistemic problems in pain management are new to the field, such as the ethical education and conduct of the computer scientists and other professionals developing AI tools in this area. Others are more perennial in nature: epistemic injustice, particularly testimonial injustice, the undermining of patient trust and self-trust, stigma, and the pervasive nature of medical

✉ Daniel Z. Buchman
daniel.buchman@utoronto.ca

[1] Institute for the History & Philosophy of Science and Technology, University of Toronto, Toronto, ON, Canada

[2] Centre for Addiction and Mental Health, Toronto, ON, Canada

[3] Department of Rhetoric and Writing, Center for Health Communication, The University of Texas at Austin, Austin, TX, USA

[4] Dalla Lana School of Public Health, University of Toronto, Toronto, ON, Canada

[5] Joint Centre for Bioethics, University of Toronto, Toronto, ON, Canada

[6] Krembil Research Institute, University Health Network, Toronto, ON, Canada

gaslighting are familiar frustrations in the pain diagnosis and treatment landscape (Bailey 2020; Fricker 2007; Jones 2012; Sebring 2021). In this paper we aim to address these concerns and highlight the ways in which proposed AI tools for pain management generate novel issues and the ways they perpetuate long standing issues related to justice — and injustice — in both medicine and AI ethics.

We first provide an overview of the current tools used to assess pain in clinical settings. We focus primarily on technologies that are used to understand the causes and qualities (e.g. intensity, severity, physical location) of a patient's pain rather than strategies used to manage pain. Here we introduce the ways that artificial intelligence (AI) is being developed for pain identification and measurement. We evaluate some of the major potential benefits and key points of concern regarding the use of these technologies. Following this overview, we present our argument, positing that there are several ethical-epistemic issues with the potential implementation of these technologies in clinical settings, drawing on literature about self-trust and epistemic and testimonial injustice. Finally, we propose that the adoption of epistemic humility on the part of both AI tool developers and clinicians can contribute to a climate of trust in and beyond pain management (Sullivan 2023) and lead to a more just approach to the implementation of AI in pain diagnosis and management. In this final section, we propose a potential methodological change to how ethical topics are discussed in the context of post-secondary education in disciplines such as medicine and computer science. At this stage, the goal of this proposal is not to suggest specific interventions but rather to promote the improvement of ethics education through an approach informed by the practice of virtue, following in the footsteps of Pellegrino and, going back much earlier, Aristotle.

Understanding pain and pain measurement today

Physicians and researchers studying pain have long understood that pain is a multidimensional health concern, and they have highlighted several models to explain the phenomenon of pain. Since Descartes, physicians have been largely interested in managing pain based on the biogenic — or biologically-based — theory of pain, wherein an underlying physical cause is responsible for a patient's discomfort (Kenny 2004, p. 297). The biogenic model of pain suggests that pain has an organic root cause that can lead to emotional secondary conditions such as depression (Peyrot, Moody, and Wiese 1993, 64). Treatment under this paradigm focuses on addressing underlying mechanisms of injury and inhibiting the transmission or reception of pain signals throughout the body. Conversely, the psychogenic model proposes that pain is a perception rather than a sensation; this means that based on the psychogenic model, pain is an expression of bodily distress (Peyrot, Moody, and Wiese 1993, 64–65). Modulation of pain response and pain behavior through cognitive-behavioral therapy is a primary mode of intervention under this model.

A third paradigm for understanding pain is the sociogenic model. This model highlights the importance of practitioners' belief in a patient's pain can demoralize or have otherwise detrimental psychological effects on a patient (Kenny 2004, p. 297). A fourth, more integrated model has been proposed to incorporate a variety of these perspectives to give a more holistic understanding of pain: the biopsychosocial model (BPSM). Initially developed by psychiatrist George Engel in 1977, the BPSM was designed to combat what Engel believed was a dogmatic and reductionist streak in the biomedical approach to medicine (Cormack et al. 2023, p. 3). Based on the theoretical framework of general systems theory, or the notion that nature, broadly speaking, is comprised of a hierarchy of levels of interacting systems, Engel proposed a view that incorporated biomedical components of disease alongside social and psychological factors (Engel 1977, p. 386). Gatchel et al. define the biopsychosocial model of understanding pain as encompassing both "objective biological disruption" of the body and the more subjective experience of how patients live with their condition (Gatchel et al. 2007, p. 582). Some see the biopsychosocial model as the gold standard for understanding pain, and pain management under this model is expressly multimodal, incorporating techniques from physical, psychological, and social care interventions (Gatchel et al. 2007, p. 607). While at first glance, this model seems to combine thoughtfully the components of biogenic, psychogenic, and sociogenic models for understanding a phenomenon like pain; however, the biopsychosocial model is not without its critics. A 2023 study by Cormack et al. argued that the way the BPSM is applied in the literature today has strayed away from Engel's original work. Citing in no small part Engel's vague introduction of his own theory, Cormack et al. argue that the BPSM is often used as Trojan Horse; the model is frequently presented with a "humanistic and holistic coating" but ultimately is applied in such a way that ignores the effects of social factors on a patient's condition (Cormack et al. 2023, p. 5).

Despite the discussion surrounding the ideal model to describe pain, most attempts to address pain outside formal pain management subspecialty contexts are vexed by a compartmentalization of pain models. Specifically, pain measurement tools currently in use in clinical settings involve some kind of scale or tool that often favour the psychogenic model of pain. These scales attempt to quantify patient perceptions of pain. For instance, a patient may be asked to rate their pain using a numerical scale from 1 to 10 (with 1 being minimal pain and 10 being the most significant amount of pain), or use a visual tool such as the Wong-Baker FACES scale (Wong-Baker FACES Foundation, n.d.). A third such tool called the Visual Analog Scale (VAS) asks patients to indicate along a line how severe their pain is (Karcioglu et al. 2018). Each of these (and other similar) tools are intended to help bridge the gap between a patient's experience of pain and a physician's understanding of their condition (see, e.g. Wiggleton-Little 2022). This psychogenic approach to pain measurement often encounters a common clinical reliance on a biogenic model of pain treatment creating a mismatch between measures and methods (Graham 2015).

Despite their limitations, psychogenic scale tools are currently among the best available to triage and help treat pain. As noted in an analysis of several of these scales, "pain cannot be treated if it is not assessed" (Karcioglu et al. 2018, p. 713). Although grounded in self-perception, these scales make room for a multicausal understanding of the experience of pain. Nevertheless, patients may find that they have trouble identifying which number best represents how they are feeling and may be concerned with whether their level of pain will accurately be conveyed to their clinician. This is especially true in pediatric populations, where pain reporting can be more complicated, and pain can go untreated or undertreated (Andersson et al. 2022; Marchetti et al. 2016). Clinicians may expect a certain degree of forthcomingness with respect to a patient making their pain level known. Pediatric patients may have a particularly difficult time relaying their phenomenological experience to their care team (Andersson et al. 2022).

This problem of phenomenology characterizes much of the challenges with these tools; while their emphasis on paring down explanatory detail can allow for faster triage (Huskisson 1974), their lack of descriptive specificity can lead to patients and clinicians "talking past" one another (Kenny 2004, p. 300). Studies have identified the range of terminology physicians and patients both use to describe their pain, and often these terms do not overlap sufficiently, as noted by the International Association for the Study of Pain (IASP) ("Terminology| International Association for the Study of Pain," n.d.). This leaves patients and clinicians frustrated and can have serious deleterious effects on the doctor-patient relationship and the efficacy of pain management (Buchman, Ho, Illes, 2016; Graham 2015; Jackson 1994; Kleinman 1994).

Trust, testimony, and the doctor-patient relationship

In recent years, philosophical literature has explored the enduring chasm between a patient's testimony and a physician's ability or willingness to effectively respond to and treat the patient's pain in a manner that respects their testimony. Much of this literature has roots in Miranda Fricker's work on epistemic and testimonial injustice, which she defines as the harm that comes from the discrediting of individuals or groups' ability to contribute to knowledge or benefit from its use and dissemination (Fricker 2007). Kidd and Carel have argued that chronically ill patients are vulnerable to multiple kinds of testimonial injustice because their being ill undermines their credibility as authorities on their own bodily condition (Kidd and Carel 2018, 216). Patients are especially vulnerable to this kind of testimonial injustice in medical settings because of the medical field's longstanding treatment of biogenic, or naturalistic definitions of health, wellbeing, illness, and pain as the objective, standard measurement of such concepts (Buchman, Ho, and Illes 2016; Kidd and Carel 2018, 216).

While vulnerability is one of the conditions many moral philosophers believe is necessary for trust (Baier 1986; Rogers and Ballantyne 2008), in time, the focus on these biogenic conceptions of these terms can erode a patient's trust in a physician, or in the healthcare system as a whole (Kenny 2004, p. 303). Buchman, Ho, and Goldberg have studied the complications of negotiating trust and epistemic injustice in the context of chronic pain management. They argue that trust is the factor that governs the emotional and interpersonal components of the doctor-patient relationship (Buchman et al. 2017, p. 32). However, as demonstrated by a large body of literature, a strong, trusting relationship is both hard to build and maintain in the context of understanding and treating chronic pain (Buchman, Ho, and Illes 2016; Kenny 2004).

Much of the issue with treating and understanding pain, especially chronic pain, comes down to communicative difficulties created by the mismatch between an individual patient's complex experience of pain and the biogenic model's attempts to reduce that complexity to a single value from a pain scale. Even though the goal is for the pain scale number to have clinical utility, it may serve an act of epistemic undermining that can reduce a patient's trust in their clinician and perhaps the medical system more broadly. In some populations, such as pediatrics, patients may be more likely to confide in a different trusted party such as a parent (Andersson et al. 2022). When a patient's complex pain experience is, in the patient's eyes, reduced to a number on a scale, their ability to trust that their care team has a full understanding of their condition may be undermined. This experience effectively triages a patient's complex condition and flattens it, meaning that even though some pain scales rely on psychogenic or even biopsychosocial models of pain, pain management often falls into a biogenic framework because the most straightforward course of action for a physician is to treat the physical pain rather than the emotional or social causes that may underlie this physical reaction. This reifies this experience of epistemic violence and the patient's subsequent experience of being distrusted.

Phenomenological and narrative approaches are often offered as alternatives to current clinical practices. For example, Havi Carel has written about the ways a patient may use phenomenology as part of their "toolkit" when interacting with healthcare providers (Carel 2012, p. 110). Graham (2015) has also explored how some clinicians even build their pain management practice around a phenomenology of suffering. Yet, in many clinical settings, there are limited opportunities for patients to engage in longform discussion with their care providers, often due to limited face-time with clinicians (E. Stahl, A. Drew, and Boer Kimball 2014). With new AI-based developments for the assessment and management of pain, this information impasse may become even more substantial potentially.

New tools

The introduction of AI in this particular corner of health care innovation is worth scrutinizing from philosophical and ethical lenses because its use in this context is tied to larger conversations on the kind of trust we ought to ascribe to AI (Alvarado 2023) and the ways that these novel technologies may alter the landscape of trust in the context of biomedicine. AI-backed tools are receiving a lot of attention for their potential to advance research and practice in pain medicine. Developers working in machine learning (ML) are particularly interested in replacing subjective patient report, even when domesticated into biogenic registers, with measurable biomarkers that have been previously indexed to subjective patient report. In so doing, they argue indicators such as biomarkers can improve pain management and contribute to more "objective, valid, and reliable" assessments of pain and diagnoses of pain-related conditions (Fernandez Rojas, Huang, and Ou 2019, 1). One recent effort, for example, aimed to develop an AI system that could reliably identify pain levels by measuring physiological responses to increasingly painful stimuli (Gruss et al. 2015). The developers of this tool argued that an "objective, reliable, and variable surrogate measurement of pain," is needed in order to replace numerical pain scales (p. 2). However, these efforts were not without a strong potential for moral hazards. Specifically, the authors indicated that their new tool would be useful for cases where patients were not "sufficiently alert and cooperative" (Gruss et al. 2015, p. 1). This invocation of cooperation implies the possibility of clinical encounters where patients can but do not consent to pain measurement or management, something that will likely provoke concern among many in bioethics and related fields. AI-driven algometry's desire for objectivity can also lead to a situation where these tools are offered as a corrective for overcoming social structural problems in pain management such as medical racism (Buchman 2024). One study algorithmically combines patient self-report of pain with clinical imaging data so as to reduce "unexplained pain disparities in underserved populations" (Pierson et al. 2021, p. 136).

This approach is quite concerning. Much of the racial inequities in pain medicine can be traced back to structural racism wherein clinicians do not trust the bodily knowledge of patients, especially when those patients belong to minority groups. As Graham (2022) has pointed out previously, "this version of "trust" does not sound a whole lot like genuine trust. If trust is only granted when verified through technological 'objectivity,' then there is no trust at all" (Graham 2022, 152). Nevertheless, proponents of incorporating ML into pain diagnosis and management tout its versatility; these tools can be used as predictors for both quantitative (diminished pain intensity) and qualitative (degrees of recovery) health outcomes (Falla et al. 2021, p. 4). These developments also allow for a wider range of potential data points. Falla et al. suggest that these tools can measure patients' pain levels between appointments in clinical settings by making use of phone applications or small wearable technologies (Falla et al. 2021, 4–5).

Another example is a 2019 study that linked several brain and autonomic markers to classify relatively higher/lower levels of pain, tracking both a patient's experience of pain over time and the comparative experiences of different patients (Lee et al. 2019). While this work is also in its early stages, Lee et al. anticipate future uses of this synergistic tool to predict pain in patients in a faster and more inexpensive manner than currently available in the clinic (Lee et al. 2019, p. 12).

This is not an isolated goal for improved diagnosis of pain. In the context of dental practice, for instance, one study found that an AI tool was able to diagnose causes of orofacial pain more accurately than dentists (Kreiner and Viloria 2022). Much like in the Lee et al. study, Kreiner and Viloria found that ML pain diagnosis tools could be used to diagnose the cause of pain faster in patients and allow for a more efficient clinical encounter (Kreiner and Viloria 2022, 888). Currently, the most applications of AI in the understanding of pain have focused on pain assessment or classification. A recent scoping review of the literature on pain AI technologies found 45 papers focused on pain diagnosis, and only 7 on clinical decision support, and four each on prognosis and self-management (Khan et al. 2024). Nagireddi et al. suggest a series of future developments that can be used to allow for more individualized treatment courses based on specific patients' medical histories. The authors posit that further development of AI in pain medicine will require more adaptable algorithms that act based on patients' individual descriptions of pain. (Nagireddi et al. 2022; E238). This would require more sophisticated ML that can more finely classify pain beyond its mere presence (Nagireddi et al. 2022; E238).

Causes for concern

One of the most touted reasons given for incorporating ML into medical practice is the potential for accuracy, and by some accounts, objectivity, though this is a more contentious claim that we will discuss further later. Numerous studies, including many of those evaluated in Nagireddi et al's study have reported that ML can accurately identify physiological causes of pain (Nagireddi et al. 2022; E238), sometimes more effectively than clinicians (Kreiner and Viloria 2022). The fact that ML does not rely on patient testimony may reassure clinicians, giving the impression of a more "objective" measure of a patient's pain level and locus/i. The use of AI in facial expression analysis is another potential instance where ML can contribute to a physician's understanding of a patient's condition without requiring patient testimony (Nagireddi et al. 2022; E236).

We can see why the development of pain diagnostic (and eventually pain management) tools that do not rely on patient testimony may offset patient burden in some contexts. Given that much of the burden patients experience in the clinic can come down to phrasing their concerns in a way that physicians both believe and understand, having a more objective measure for the location and level of their pain may offset a patient's need to advocate for themselves to the same extent as without an AI-backed diagnosis. Furthermore, patients who are unable to communicate with their care teams (either due to language barriers or conditions that limit or otherwise inhibit speech) may benefit from the implementation of an AI-powered diagnostic tool as their ability to be treated is no longer reliant on their ability to relay their experience to physicians.

The introduction of tools that attempt to verify patient testimony can also potentially bolster a patient's intellectual self-trust, a phenomenon Karen Jones describes as "an attitude of optimism about one's cognitive competence in a domain" (Jones 2012, p. 245). An individual with a high degree of self-trust in a particular area, such as their pain experience, will be likely to have strong convictions of their ability or knowledge in that area (Jones 2012, p. 243). Given that historically, belief in a patient's pain has been a point vulnerable to communication breakdown in clinical settings (Kenny 2004), the use of ML tools may corroborate a patient's experience and empower them by proving their experience to a physician.

Despite the considerable optimism that these tools can make parts of the diagnostic process easier for both patients and physicians, they are not without their potential drawbacks. For instance, while these technologies may help bolster patient self-trust and empower them to feel assured in relaying their experiences, a result that contradicts the patient's self-reported pain level may come as a blow to the patient's self-trust. In the context of pain diagnosis, this degradation of self-trust may be tied to what Miranda Fricker calls hermeneutical injustice, where an individual is left without the resources to understand, and therefore relay their experiences (Fricker 2007). This experience can lead to what Jones refers to as the corrosion of self-trust (Jones 2012, p. 246).

Scholars are already warning of the potential — or reality — for new forms of hermeneutical injustice to occur with the advent of new ML technologies in other health care spaces. Hermeneutical injustice can already be seen in the application of algorithms used to monitor and predict patients' potential to misuse opioids (Pozzi 2023, p. 4). These algorithms can "misunderstand" basic, innocuous facts, such as consulting a second doctor for an opinion or traveling from a rural area to a larger city or town for medication as risk factors for opioid misuse or addiction, and since these algorithms are in an unwarranted but nonetheless privileged epistemic position, their "word" is taken at the expense of a patient's experience (Pozzi 2023, p. 5). Another such case is conversational AI (CAI), such as chatbots. CAI have limited "vocabularies" for patients to use to express and understand their own condition, and if a patient's experience falls outside of what that vocabulary can express, the patient's condition will go unacknowledged (de Proost and Pozzi 2023, p. 52). In both these and other cases, the outsourcing of medical work to ML tools can degrade both a patient's self-trust in their own understanding of their symptoms and reduce their trust in their doctor(s), or the medical system more generally. On the other side of the doctor-patient relationship, ML tools in the world of pain diagnosis may contribute to an overemphasis on self-trust from the clinician's perspective, and over time this may distort their understanding of the limits of their own expertise (Jones 2012, p. 246).

ML tools can also contribute to the supposed objectivity of medical treatment, especially in the world of pain medicine. Objective data are generally deemed to be epistemically superior to more individualistic evidence, such as patient testimony (Buchman et al. 2017). Anita Ho has discussed this further, arguing that objectivity is often ascribed to tools such as algorithms and that this can further the effects of ongoing biases in medicine (Ho 2022, p. 295), such as diminished belief in patient testimony. Ho fears that the authority both the public and physicians bestow upon AI tools can lead to or worsen the types of epistemic injustice that patients may experience when interacting with the medical system (Ho 2022, p. 295). We discuss this issue in greater detail in the following section.

Other causes for concern relate to the availability and quality of the data gathered by these tools. In their 2018 overview paper of the state of machine learning in pain research, Lötsch and Ultsch highlight a few of these shortcomings, most notably that algorithms can only capture data they have already been trained to recognize (Lötsch and Ultsch 2018, 628). This means that algorithms may miss data that is "unfamiliar" to them. Lötsch, Ultsch, and coauthors elaborate on this in a more recent paper, concluding more strongly that AI and ML tools in pain research can expedite repetitive identification work, but that ideally both physicians and technological tools can be used to yield an ideal result for patients at this time (Lötsch et al. 2022). The technical and epistemic shortcomings of these tools, combined with the issues of self-trust contribute to the furthering of a clinical environment that overemphasizes the authority and expertise of clinicians while undermining patient experience and testimony.

Part of the issue with these tools is due to the way they are (mis)used or unable to be used in some clinical settings. In *The Doctor and the Algorithm*, Graham highlights that these tools may not serve a physician trying to gauge an unconscious — or non-compliant — patient's pain level (Graham 2022, 145–46). Tangled in with this is the issue of physicians' belief in a patient's pain, especially in the context of chronic pain (Kenny 2004). Both Graham and Kenny highlight that doctors look for objective or biogenic causes for pain, and diagnosing long-lasting, non-specific pain that seemingly has no physical root cause is both difficult and requires the mutual placing of trust to encourage belief in either the pain (on the physician's part) or in the physician's belief in the patient's pain (Graham 2022; Kenny 2004).

Epistemic-ethical burdens for patients

There are a number of epistemic and ethical issues that come to the fore with the presence of machine learning and other AI-powered tools in clinical settings, mainly as additional burdens placed on ill patients who are already vulnerable to a range of injustices (Kidd and Carel 2018, 219).

Believing patients

One of the most pressing shortcomings of the use of AI in the clinic is the potential additional burden these tools may place on a patient. This burden has two components: (1) testimonial deficit and (2) model patient behaviour. Suppose a patient sees a doctor for a high intensity persistent pain in their lower back, which is not relieved by analgesics such as ibuprofen or acetaminophen. The doctor refers the patient for a few tests, including a new test that looks for biomarkers that may identify the cause of the patient's pain. This test yields no conclusive results, and, heavily relying on the lack of an answer given by this test, the physician dismisses the patient's concern and recommends they continue taking over-the-counter analgesics until the pain is resolved. This experience downplays the patient's own experience, and while they may be able to make use of the kinds of tools Carel proposes in her arguments for the use of phenomenology as a patient resource (Carel 2012), it is still likely that this experience may leave the patient at a testimonial deficit. The predisposition to testimonial injustice in the clinic means that patients already at risk of encountering testimonial injustice are not only likely to be subjected to such injustice, but they will be forced to testify *against* a tool touted as an "objective" measure for pain that a physician may be more willing to use as the basis for treatment decisions than a patient's testimony.

This experience can be classified in several ways, none of them positive. A patient may feel angry or frustrated, or taken advantage of, or dismissed by an authority figure whose job is to care for them when they are unwell. This can lead to patients presenting as uncooperative or combative, and in more extreme cases, perhaps even "non-compliant". Non-compliance is a serious barrier to treatment, but it can be detrimental to use this kind of terminology to label patients who find they must be assertive or even combative at times to have their testimony taken seriously alongside AI-powered test results combined with physician expertise. Often, patients advocate for themselves in this capacity in order to play an active role in their own treatment decisions rather than to be passive recipients of care (Playle and Keeley 1998). This "combativeness" is often required from certain groups of patients, such as racialized patients or those who use drugs, to be taken seriously. Graham has expressed further concern over the use of the terms "non-compliant" and "sufficiently cooperative", highlighting the constraints the connotations of these terms may place on patient autonomy (Graham 2022, 146). Often for these groups, if they are too expressive then they are pain catastrophizing (another stigmatized term), but if they minimize their responses to avoid the catastrophizing label, they are considered not in enough pain.

Transparency in clinical practice

There have long been calls for transparency in the development of common medical practice. A few decades prior, the supremacy of evidence-based medicine (EBM) was a popular debate in philosophy of medicine. As succinctly defined by Sackett et al. in 1996, EBM "is the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients" (Sackett et al. 1996, p. 71). Favoured for its bottom-up philosophy, EBM is often lauded for its ability to combine clinical observation with broader expertise and patient choice (Sackett et al. 1996, p. 72). Ideally, EBM is practiced by incorporating clinical experience and expertise with rigorous research external to a particular clinical setting (Sackett et al. 1996, p. 71).

This is not always the case though, and despite the praise EBM has received, it has also been subject to criticism, some of which can be seen repeated in the current discourse about the use of AI in medical practice. EBM has been critiqued for its assumption of value neutrality in science and in medicine and its reinforcement of physicians' authority in medical settings (Goldenberg 2006). EBM attempts to separate the sociopolitical elements of health care (e.g. a patient's demographic) from the technical elements (e.g. a particular test of procedure), an approach that can limit patient input regarding their own care (Goldenberg 2006, 2630). One of the strongest critiques of EBM describes the methodology as a "black box", arguing that the arguments and decisions behind its mechanisms have become obscured over time (Valkenburg, Achterhuis, and Nijhof 2003, 468). In this way, EBM is just one other example of the longstanding evaluative tools and methods that can, even unintentionally, discredit patient experiences and reduce patient trust. In more recent years, an amended version of EBM, referred to EBM+has been proposed. Proponents of EBM+argue that this "new and improved" model for practicing medicine additionally places emphasis on the underlying mechanism(s) that make an intervention work for a particular course of treatment (Pfaff and Schmitt 2023). While EBM+has been proposed to account for the previously underexplored elements of care in medicine, issues remain with the degree of transparency patients can expect about their care or course of treatment.

Algorithmic opacity is a cause for concern in healthcare and many other fields. While these tools are increasingly reliable, sometimes even more so than human experts, their "reasoning" is unknown to the human experts using these tools (von Eschenbach 2021, 1614). To a patient, especially a pain patient already disenfranchised by the kinds of testimonial injustice they have likely encountered in the clinic, a diagnostic tool powered by an algorithm can reify the technical knowledge a clinician can access but a patient cannot without the assistance of an expert. This could cause a patient to doubt their own experience of pain and reduce their own self-trust (Jones 2012, p. 246).

Offsetting patient burden

The problems laid out in the previous sections of this paper are diverse in nature, and no singular solution will act as a panacea. However, at this point we propose that the adoption of an approach to both developing and using AI tools in pain medicine based in epistemic humility may allow for more fruitful discussion and treatment in the clinic — both with respect to a patient's condition and their overall treatment as a person — and contribute to a more just clinical environment in pain management.

Epistemic humility

José Medina argues that epistemic humility allows one to adopt a "self-questioning attitude toward one's cognitive repertoire" (Medina 2013, p. 43). This humility-based approach to understanding provides an environment where one can then identify gaps in one's knowledge or expertise and consequently contribute to one's desire to learn the material required to fill those gaps (Medina 2013, p. 43). Epistemic humility requires that people — especially those in positions of power — internally examine and critique their own claims to cognitive authority in order to ensure that other perspectives are not only heard but incorporated into furthering understanding (Potter 2002, p. 123).

On the surface, this seems like a simple solution to the problem of unfairly undermining the credibility of patients in the context of pain management. Indeed, there are calls for epistemic humility in AI development and AI ethics more broadly (Gellers 2023; McCradden, Hui, and Buchman 2023). Most clinicians would probably claim that they thoughtfully listen to patient concerns, have their patient's interests as primary, and are epistemically humble in their approach. We share that assumption. However, we know from the extensive humanities and social science literature that this experience is not shared by many people in pain, and the potential introduction of AI tools in this context creates the conditions for intensifying epistemic injustice among those most vulnerable to it. Accordingly, the adoption of an epistemically humble approach must be developed and enacted in a careful and particular manner to ensure that this move is true to the spirit of the concept of epistemic humility rather than simply serving as lip service.

How is such an approach achieved? We propose that much of the solution lies in earlier and more prominent ethics education in non-philosophy disciplines. For this paper, we turn to two disciplines: medicine and engineering. The former's inclusion is straightforward; future doctors ought to be equipped with the moral skillset to not only run their clinics in an ethical manner, but to make use of Medina's "self-questioning" approach in their work (Medina 2013,

p. 43). The inclusion of engineering is intended to stress the importance of stronger ethics education in the discipline that often trains the programmers and developers who design and build ML and other AI tools for use in clinical settings. This call to physicians, engineers, and other professionals in the healthcare sector has been made before in reports examining the current and potential future uses of AI in healthcare (Matheny et al. 2019, p. 3).

Ethics education

While ethics education is considered to be an important part of medical education (Stites et al. 2020), ethics — and humanities education more broadly — is often underexplored and sacrificed at the expense of technical medical and scientific curriculum material (Kollmer Horton 2019). Students often receive insufficient training on the ethical issues present in their chosen profession, and the addition of a single course on these issues is unrealistic given a variety of constraints on these programs already (Shelton 1999, p. 673).

As with many papers on topics related to ethical conduct in medical practice and innovation, we advocate strongly for more comprehensive ethics education for engineering and medical trainees. Much of the existing work on implementing ethics education into engineering curricula is narrow in scope and largely revolves around structured ethical models (e.g. the Trolley Problem), or extreme case studies that imply that everyday ethics issues are not commonly encountered in engineering (Furey and Martin 2019; Pierrakos et al. 2019). These models often lack the kind of context that is applicable to medical or engineering students. Instead of advocating for specific curricular interventions for medical students, biomedical engineers, and other professional groups, we want to encourage a different approach. Newer projects, such as “embedded ethics” education initiatives (sometimes stylized as Embedded EthiCS or E3I depending on the institution) seek to meet engineering and computer science students in their home classrooms. These modules consist of lectures and activities on pressing topics related to ethics of technology and AI, including ethical issues in the creation of addictive algorithms, data collection and privacy, and responsibility within AI innovation (“E3I: Embedded Ethics Education Initiative,” n.d.; “Embedded EthiCS @ Harvard,” n.d.).

However, as researchers who work parallel to and within the medical field, we find these calls for improved ethics seminars and courses are often too vague to be practical. In advocating for an improvement to medical and engineering ethics education, we want to suggest an approach that encourages students (and educators) to think of ethics education as being holistic approach to one’s decision-making rather than a simple professional checklist or obligation. As argued by previous advocates in this space, most ethics educators in these academic contexts are “teaching *about* ethics”, a pedagogical detail that may allow students to pass a test or course, but will not alter their behaviour or thought processes pertaining to long-term care, treatment, and professional development (Bertolami 2004, p. 414).

What could this pedagogical reframing resemble? One option is to embrace the Aristotelian approach to teaching ethical conduct much in the same way that other kinds of virtues can be taught. While other forms of ethical analysis could be employed well in this area, virtue-based approaches may be particularly useful in the context of teaching medical and engineering students. Instructors and mentors in these educational contexts typically have professional experience in the fields they teach in, and they are in the somewhat unique position of teaching their students both by way of formal educational lessons and by demonstrating and embodying the virtues they wish to instill in their students.

Bioethicist Edmund Pellegrino advocated for a virtue ethics turn in the context of medical education, arguing that in part, ethical physicians develop their virtues by observing and emulating their mentors (Pellegrino 2002, p. 383). Aristotle’s conception of virtue can be summarized as the “excellence in function”, a set of lessons learned and internalized, not simply memorized and regurgitated (Shelton 1999, p. 672). Shelton argues that the primary goal of ethics educators in medical schools (and presumably other professional programs) is to “embody and exemplify the habits of virtue in action” (Shelton 1999, p. 674). This view is echoed by Pellegrino, who argues that a great deal of character formation can still occur when a student is in medical school and is exposed to clinical instructors who exhibit the virtues that strong ethics education ought to discuss and demonstrate (Pellegrino 2002, p. 383), including epistemic humility.

However not all medical ethicists have historically supported this perspective. Robert Veatch argued that it is misguided to make assumptions about there being a singular set of virtues that all “social roles” (in this case, would-be engineers and doctors) would benefit from embodying (Veatch 1985). Veatch adds that because there is no set standard virtues, it is possible to condone morally wrong acts using a virtue ethics framework (Veatch 1985, p. 334). These arguments extend beyond the applied realm of bioethics. Martha Nussbaum has similarly criticized the spotlight virtue ethics shares with utilitarianism and Kantianism as a major ethical system taught in introductory ethics classrooms (Nussbaum 1999). Nussbaum argues that unlike Kantian or utilitarian ethics, virtue ethics cannot be readily digested into a neat code or set of rules or steps for one to follow (Nussbaum 1999, p. 167). While these critiques of virtue ethics have their merits, one of the advantages of delivering ethics instruction through a virtue ethics approach is precisely what Nussbaum criticizes. Using a primarily Kantian or utilitarian approach may lead to students and trainees assuming that all ethical dilemmas can be resolved by way of a simple moral calculus.

Under this view, virtue and ethics education moves away from resembling a checklist or even a simple, rigid code of conduct. In the context of dental ethics education, for example, Bertolami advocates for an ethics education model that presents real opportunities for growth, not just in one's professional life, but also in their personal life (Bertolami 2004, p. 422). This is a position we can also defend. Bertolami pushes dental schools to move to an ethics education model where more general moral introspection is encouraged without fear of humiliation. He writes that schools must create "a climate that helps students understand that it isn't enough just to do the right thing, that what is also required is to *want* to do the right thing" (Bertolami 2004, p. 423). This sense of *desire* for an ethical resolution, much like Shelton's call for ethics educators to exemplify the *habit* of virtue, would help ethics education transform from being perceived as a single mark on a transcript to an essential component of both personal conduct and professional practice.

A key question remains here: would an overhaul of the current ethics curriculum in medical and engineering programs *teach* epistemic humility? That depends on what you mean by 'teach.' Medical education famously integrates didactic and practical instructional modes finding that effective care if best supported when students have supervised opportunities to deploy technical knowledge in situated practice. Thus, perhaps the *practice* of character growth can be taught in a similar vein to many other skills students learn (Pierrakos et al. 2019, 5–6), i.e., through a curriculum that combines didactic and practical ethical education. Additionally, students can be taught the benefits of bolstering self-trust in peers (where appropriate) and in populations that stand to be affected by their interventions (i.e. medical care or tool development). In teaching students these two skills in tandem with one another, they will likely come to see the limits of their own expertise and understand where they ought to consider or rely upon the expertise and testimony of others, whether they be experts in their own fields or patients with the lived experience of their own pain. This approach to framing how ethics is taught, in tandem with the demonstration by mentors of what it means to be a virtuous, ethical, professional may reinvigorate ethics curricula and improve patients' experiences with both human practitioners and cutting-edge AI technologies within the medical system.

Conclusion

When faced with the rapid development of new technology, it is often difficult for those who work parallel to health research to determine where these new technologies can benefit — or hinder — patient care. In this paper, we have drawn attention to just one area in medicine where this debate is particularly fraught. The use of AI tools in diagnosing pain and pain conditions is currently quite limited, but as this technology continues to grow it will no doubt be used in a wider array of contexts within pain management. If these tools are implemented without sufficient care, the healthcare system risks worsening the trust of a population that is already prone to experiencing stigma and epistemic injustice in clinical settings. As researchers from diverse backgrounds, we hope to see strategic, powerful, and effective developments in the instruction medical and engineering students receive concerning epistemic humility and authority in their professional practice.

Acknowledgments The authors wish to thank Dr Ross Upshur for helpful suggestions of literature on ethics education.

Funding RAK reports funding from the Social Sciences and Humanities Research Council of Canada; SSG receives grant support from the National Institute of General Medical Sciences of the National Institutes of Health, unrelated to the current project; DZB reports funding from the Canadian Institutes of Health Research Project Grant (PJT – 178023).

Declarations

Conflict of interest The authors report no conflicts of interest.

References

Alvarado, Ramón, and If Any?. 2023. *AI and Ethics* 3 (4): 1169–1183. https://doi.org/10.1007/s43681-022-00224-x.

Andersson, Viveka, Stefan Bergman, Ingela Henoch, and Hanna Simonsson, Karin Ahlberg. 2022. Pain and pain management in children and adolescents receiving hospital care: A

Cross-Sectional study from Sweden. *BMC Pediatrics* 22(1):252. https://doi.org/10.1186/s12887-022-03319-w

Baier, Annette C. 1986. Trust and antitrust. *Ethics* 96(2):231–260.

Bailey, Alison. 2020. On gaslighting and epistemic injustice: Editor's introduction. *Hypatia* 35(4):667–673. https://doi.org/10.1017/hyp.2020.42

Bertolami, Charles N. 2004. Why our ethics curricula don't work. *Journal of Dental Education* 68(4):414–425. https://doi.org/10.1002/j.0022-0337.2004.68.4.tb03758.x

Buchman, Daniel Z. 2024. AI and the ethics of Techno-Solutionism in pain management. *PAIN*. https://doi.org/10.1097/j.pain.0000000000003389

Buchman, Daniel Z., and Anita Ho, Judy Illes. 2016. You present like a drug addict: Patient and clinician perspectives on trust and trustworthiness in chronic pain management. *Pain Medicine* 17(8):1394–1406. https://doi.org/10.1093/pm/pnv083

Buchman, Daniel Z., Anita Ho, and Daniel S. Goldberg. 2017. Investigating trust, expertise, and epistemic injustice in chronic pain. *Journal of Bioethical Inquiry* 14(1):31–42. https://doi.org/10.1007/s11673-016-9761-x

Carel, Havi. 2012. Phenomenology as a resource for patients. *The Journal of Medicine and Philosophy* 37(2):96–113. https://doi.org/10.1093/jmp/jhs008

Cormack, Ben, Peter Stilwell, and Sabrina Coninx, Jo Gibson. 2023. The biopsychosocial model is lost in translation: From misrepresentation to an enactive modernization. *Physiotherapy Theory and Practice* 39(11):2273–2288. https://doi.org/10.1080/09593985.2022.2080130

E3I: Embedded Ethics Education Initiative. n.d. Accessed October 1. 2024. https://www.cs.toronto.edu/embedded-ethics/

Embedded EthiCS @ Harvard. n.d. Accessed October 1. 2024. https://embeddedethics.seas.harvard.edu/

Engel, George L. 1977. The need for a new medical model: A challenge for biomedicine. *Psychodynamic Psychiatry* 40(3):377–396. https://doi.org/10.1521/pdps.2012.40.3.377

Falla, Deborah, Valter Devecchi, David Jiménez-Grande, David Rügamer, and Bernard X. W. Liew. 2021. Machine learning approaches applied in spinal pain research. *Journal of Electromyography and Kinesiology* 61(December):102599. https://doi.org/10.1016/j.jelekin.2021.102599

Fernandez Rojas, Raul, Xu, Huang, and Keng-Liang Ou. 2019. A machine learning approach for the identification of a biomarker of human pain using fNIRS. *Scientific Reports* 9(1):5645. https://doi.org/10.1038/s41598-019-42098-w

Fricker, Miranda. 2007. *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.

Furey, Heidi, Fred Martin. 2019. AI education matters: A modular approach to AI ethics education. *AI Matters* 4(4):13–15. https://doi.org/10.1145/3299758.3299764

Gatchel, Robert J., Yuan Bo Peng, Madelon L. Peters, N. Perry, and Fuchs, Dennis C. Turk. 2007. The biopsychosocial approach to chronic pain: Scientific advances and future directions. *Psychological Bulletin* 133(4):581–624. https://doi.org/10.1037/0033-2909.133.4.581

Gellers, Joshua C. 2023. AI Ethics Discourse: A Call to Embrace Complexity, Interdisciplinarity, and Epistemic Humility. *AI & SOCIETY*, June. https://doi.org/10.1007/s00146-023-01708-y

Goldenberg, Maya J. 2006. On Evidence and Evidence-Based Medicine: Lessons from the Philosophy of Science. *Social Science &* Medicine, Part Special Issue: Gift Horse or Trojan Horse? Social Science Perspectives on Evidence-based Health Care, 62 (11): 2621–32. https://doi.org/10.1016/j.socscimed.2005.11.031

Graham, S., and Scott. 2015. *The politics of pain medicine: A Rhetorical-Ontological inquiry*. Chicago: The University of Chicago Press.

Graham, S., and Scott. 2022. Ethics, justice, and health AI. In *The Doctor and the algorithm: Promise, peril, and the future of health AI*, 264. Oxford University Press.

Gruss, Sascha, Roi Treister, Philipp Werner, Harald C. Traue, Stephen Crawcour, and Adriano Andrade, Steffen Walter. 2015. Pain intensity recognition rates via biopotential feature patterns with support vector machines. *PLOS ONE* 10(10):e0140330. https://doi.org/10.1371/journal.pone.0140330

Ho, Anita. 2022. Artificial intelligence as a feminist bioethics issue. In *The Routledge handbook of feminist bioethics*, Routledge.

Huskisson, E. C. 1974. Measurement of pain. *The Lancet*, originally published as 2, Issue 7889, 304 (7889): 1127–1131. https://doi.org/10.1016/S0140-6736(74)90884-8

Jackson, Jean E. 1994. 'After a While No One Believes You': Real and Unreal Pain. In *Pain as Human Experience*, edited by Mary-Jo DelVecchio Good, Paul Brodwin, Byron J. Good, and Arthur Kleinman, 138–68. University of California Press. https://doi.org/10.1525/9780520354722-007

Jones, Karen. 2012. The politics of intellectual Self-Trust. *Social Epistemology* 26(2):237–251. https://doi.org/10.1080/02691728.2011.652215

Karcioglu, Ozgur, Hakan Topacoglu, Ozgur Dikme, and Ozlem Dikme. 2018. A systematic review of the pain scales in adults: Which to use?? *The American Journal of Emergency Medicine* 36(4):707–714. https://doi.org/10.1016/j.ajem.2018.01.008

Kenny, Dianna T. 2004. Constructions of Chronic Pain in Doctor–Patient Relationships: Bridging the Communication Chasm. *Patient Education and Counseling*, Special Section: Chronic Obstructive Pulmonary Disease, 52 (3): 297–305. https://doi.org/10.1016/S0738-3991(03)00105-8

Khan, Md, and Asif et al. 2024. Cracking the chronic pain code: A scoping review of artificial intelligence in chronic pain research. *Artificial Intelligence in Medicine* 151:102849.

Kidd, Ian James. 2018. and Havi Carel. Healthcare Practice, Epistemic Injustice, and Naturalism. In *Harms and Wrongs in Epistemic Practice*, edited by Simon Barker, Charlie Crerar, and Trystan S. Goetze. Wellcome Trust–Funded Monographs and Book Chapters. Cambridge (UK): Cambridge University Press. http://www.ncbi.nlm.nih.gov/books/NBK562587/

Kleinman, Arthur. 1994. Pain and Resistance: The Delegitimation and Relegitimation of Local Worlds. In *Pain as Human Experience*, edited by Mary-Jo DelVecchio Good, Paul Brodwin, Byron J. Good, and Arthur Kleinman, 169–97. University of California Press. https://doi.org/10.1525/9780520354722-008

Kollmer Horton, Mary E. 2019. The orphan child: Humanities in modern medical education. *Philosophy Ethics and Humanities in Medicine* 14(1):1. https://doi.org/10.1186/s13010-018-0067-y

Kreiner, Marcelo, Jesús, and Viloria. 2022. A novel artificial neural network for the diagnosis of orofacial pain and temporomandibular disorders. *Journal of Oral Rehabilitation* 49(9):884–889. https://doi.org/10.1111/joor.13350

Lee, Jeungchan, Ishtiaq Mawla, Jieun Kim, Marco L. Loggia, Ana Ortiz, Changjin Jung, and Suk-Tak Chan et al. 2019. Machine Learning-Based prediction of clinical pain using multimodal neuroimaging and autonomic metrics. *Pain* 160(3):1–23. https://doi.org/10.1097/j.pain.0000000000001417

Lötsch, Jörn, Alfred Ultsch. 2018. Machine learning in pain research. *Pain* 159(4):623–630. https://doi.org/10.1097/j.pain.0000000000001118

Lötsch, Jörn, Alfred Ultsch, and Benjamin Mayer, Dario Kringel. 2022. Artificial intelligence and machine learning in pain research: A data scientometric analysis. *Pain Reports* 7(6):e1044. https://doi.org/10.1097/PR9.0000000000001044

Marchetti, G., A. Vittori, V. Tortora, M. Bishop, G. Lofino, V. Pardi, and E. A. De Marco et al. 2016. Prevalence of pain in the departments of surgery and oncohematology of a paediatric hospital that has joined the project 'towards pain free hospital'. *La Clinica Terapeutica* 167(5):156–160. https://doi.org/10.7417/CT.2016.1948

Matheny, Michael, Sonoo Thadaney Israni, and Mahnoor Ahmed, Danielle Whicher. 2019. *Artificial intelligence in health care: The hope, the hype, the promise, the peril. 10*. Washington DC: National Academy of Medicine.

Mathur, Vani A., Zina Trost, Miriam O. Ezenwa, and John A. Sturgeon, Anna M. Hood. 2022. Mechanisms of injustice: What we (Do Not) know about Racialized disparities in pain. *PAIN* 163(6):999. https://doi.org/10.1097/j.pain.0000000000002528

McCradden, Melissa, Katrina Hui, Z. Daniel, and Buchman. 2023. Evidence, ethics and the promise of artificial intelligence in psychiatry. *Journal of Medical Ethics* 49(8):573–579. https://doi.org/10.1136/jme-2022-108447

Medina, José. 2013. Resistance as Epistemic Vice and as Epistemic Virtue. In *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination*, edited by José Medina, 0. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199929023.003.0002

Nagireddi, Jagadesh N, Amanya Ketan Vyas, R. Mahendra, Amol Sanapati, and Soin. 2022. and Laxmaiah Manchikanti. The Analysis of Pain Research through the Lens of Artificial Intelligence and Machine Learning. *Pain Physician*.

Nussbaum, Martha C. 1999. Virtue ethics: A misleading category?? *The Journal of Ethics* 3(3):163–201. https://doi.org/10.1023/A:1009877217694

Pellegrino, Edmund D. 2002. Professionalism, profession and the virtues of the good physician 69 (6).

Peyrot, Mark, and Philip M. Moody, H. Jean Wiese. 1993. Biogenic, psychogenic, and Sociogenic models of adjustment to chronic pain: An exploratory study. *The International Journal of Psychiatry in Medicine* 23(1):63–80. https://doi.org/10.2190/PFKJ-7NNU-WVLH-33W7

Pfaff, Holger, Jochen Schmitt. 2023. Reducing uncertainty in Evidence-Based health policy by integrating empirical and theoretical evidence: An EbM+theory approach. *Journal of Evaluation in Clinical Practice* 29(8):1279–1293. https://doi.org/10.1111/jep.13890

Pierrakos, Olga, Mike Prentice, Cameron Silverglate, Michael Lamb, and Alana Demaske. 2019. and Ryan Smout. Reimagining Engineering Ethics: From Ethics Education to Character Education. In *2019 IEEE Frontiers in Education Conference (FIE)*, 1–9. https://doi.org/10.1109/FIE43999.2019.9028690

Pierson, Emma, David M. Cutler, Jure Leskovec, and Sendhil Mullainathan, Ziad Obermeyer. 2021. An algorithmic approach to reducing unexplained pain disparities in underserved populations. *Nature Medicine* 27(1):136–140. https://doi.org/10.1038/s41591-020-01192-7

Playle, John F., Philip Keeley. 1998. Non-Compliance and professional power. *Journal of Advanced Nursing* 27(2):304–311. https://doi.org/10.1046/j.1365-2648.1998.00530.x

Potter, Nancy Nyquist. 2002. *How can I be trusted?? A virtue theory of trustworthiness*. Rowman & Littlefield.

Pozzi, Giorgia. 2023. Automated opioid risk scores: A case for machine Learning-Induced epistemic injustice in healthcare. *Ethics and Information Technology* 25(1):3. https://doi.org/10.1007/s10676-023-09676-z

de Proost, Michiel, Giorgia Pozzi. 2023. Conversational artificial intelligence and the potential for epistemic injustice. *The American Journal of Bioethics* 23(5):51–53. https://doi.org/10.1080/15265161.2023.2191020

Rogers, Wendy, Angela Ballantyne. 2008. Gender and trust in medicine: Vulnerabilities, abuses, and remedies. *IJFAB: International Journal of Feminist Approaches To Bioethics* 1(1):48–66. https://doi.org/10.3138/ijfab.1.1.48

Sackett, D. L., W. M. Rosenberg, J. A. Gray, R. B. Haynes, and W. S. Richardson. 1996. Evidence based medicine: what it is and what it isn't. *BMJ: British Medical Journal* 312(7023):71–72.

Sebring, Jennifer C. H. 2021. Towards a sociological Understanding of medical gaslighting in Western health care. *Sociology of Health & Illness* 43(9):1951–1964. https://doi.org/10.1111/1467-9566.13367

Seers, Tim, Sheena Derry, Kate Seers, and R. Andrew Moore. 2018. Professionals underestimate patients' pain: A comprehensive review. *PAIN* 159(5):811. https://doi.org/10.1097/j.pain.0000000000001165

Shelton, W. 1999. Can virtue be taught?? *Academic Medicine* 74(6):671.

Stahl, E., Mark A. James, and Drew, Alexa Boer Kimball. 2014. Patient-Clinician concordance, Face-Time and access. *International Journal of Health Care Quality Assurance* 27(8):664–671. https://doi.org/10.1108/IJHCQA-04-2013-0048

Stites, S. D., S. Rodriguez, C. Dudley, and A. Fiester. 2020. Medical students' exposure to ethics conflicts in clinical training: Implications for timing UME bioethics education. *Hec Forum* 32(2):85–97. https://doi.org/10.1007/s10730-020-09412-w

Sullivan, Laura Specker. 2023. Climates of distrust in medicine. *Hastings Center Report* 53(S2):S33–38. https://doi.org/10.1002/hast.1521

Terminology| International Association for the Study of Pain. n.d. International Association for the Study of Pain (IASP) (blog). Accessed December 17. 2021. https://www.iasp-pain.org/resources/terminology/

Valkenburg, Govert, Hans Achterhuis, André, and Nijhof. 2003. Fundamental shortcomings of Evidence-based medicine. *Journal of Health Organization and Management* 17(6):463–471. https://doi.org/10.1108/14777260310506614

Veatch, Robert M. 1985. Against Virtue: A Deontological Critique of Virtue Theory in Medical Ethics. In *Virtue and Medicine: Explorations in the Character of Medicine*, edited by Earl E. Shelp, 329–45. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-009-5229-4_17

von Eschenbach, Warren J. 2021. Transparency and the black box problem: Why we do not trust AI. *Philosophy & Technology* 34(4):1607–1622. https://doi.org/10.1007/s13347-021-00477-0

Wiggleton-Little, Jada. 2022. MPQ descriptors as a pain pidgin Language. *Journal of Hospital Ethics* 8(2):85–89.

Wong-Baker, FACES Foundation. n.d. Home. Wong-Baker FACES Foundation. Accessed December 18. https://wongbakerfaces.org/