Originalartikel

Between Theoretical and Empirical Ethics of Healthcare Artificial Intelligence

The Case of Autonomy in Breast Cancer Screening

Victor Vadmand Jensen[1,2]
Jens Christian Bjerring[3]

[1]Institut for Klinisk Medicin, Aarhus Universitet
[2]Universitetsklinik for Interdisciplinære Ortopædkirurgiske Forløb, Regionshospitalet Silkeborg
[3]Institut for Kultur og Samfund, Aarhus Universitet
vvj@clin.au.dk
filjcb@cas.au.dk

Jensen, Victor Vadmand & Bjerring, Jens Christian. 2026. Between Theoretical and Empirical Ethics of Healthcare Artificial Intelligence: The Case of Autonomy in Breast Cancer Screening. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 124-147. DOI: 10.7146/tfss.v25i44.155084

Indsendt 05/25, accepteret 03/26, udgivet 06/26

The ethical challenges of healthcare artificial intelligence (AI) have received widespread attention, as they can undermine trust and limit AI's potential benefits. A key concern is the delegation of autonomy to AI systems, a practice that has drawn criticism from practitioners, policymakers, and researchers. While theoretical ethical guidelines emphasize restricting AI autonomy, they are often seen as vague or impractical. In response, scholars have proposed an empirical ethics approach, which grounds ethical considerations in real-

world clinical settings. However, little research has examined how this approach applies to healthcare AI in practice. This paper contrasts theoretical and empirical ethics in the context of AI autonomy. Using breast cancer screening in Danish healthcare as a case study, we explore how intra-normativity shapes perceptions of good care. Our findings show that while theoretical discourse deems autonomous AI unethical, clinical practice views it as ethically acceptable. We conclude by discussing how empirical ethics can inform more practical, context-sensitive guidelines for healthcare AI implementation.

Imellem Teoretisk og Empirisk Etik for Kunstig Intelligens i Sundhedsvæsenet: En Case om Autonomi i Brystkræftscreening

De etiske udfordringer forbundet med kunstig intelligens (AI) i sundhedsvæsenet har fået stor opmærksomhed, da de kan underminere tillid og begrænse AI's potentielle fordele for sundhedsvæsenet. En særlig etisk udfordring er delegeringen af autonomi til AI-systemer, som er blevet fordømt af praktikere, politikere og forskere. Selvom teoretiske etiske guidelines har fokuseret på at begrænse AI's autonomi, er de blevet kritiseret for at være vage og upraktiske. Forskere har responderet med forslag om at bruge empirisk etik, hvor etik kontekstualiseres og situeres i klinisk praksis. Der er dog begrænset forskning om, hvordan empirisk etik praktisk kan anvendes til AI i sundhedsvæsenet. I denne artikel sidestiller vi teoretisk og empirisk etik i konteksten af AI-systemers autonomi. Med brystkræftscreening som case undersøger vi, hvordan intranormativitet former særegne forestillinger om god omsorg. Vores resultater viser, at selvom teoretisk diskurs ser autonom AI som uetisk, så ser klinisk praksis det som etisk acceptabelt. Vi slutter vores artikel med at diskutere, hvordan vi kan bruge empirisk etik om AI i sundhedsvæsenet til at skabe praktiske og kontekst-sensitive guidelines for AI i sundhedsvæsenet.

Introduction

We are witnessing rapid development and implementation of artificial intelligence (AI) systems in healthcare. These systems are envisioned to support clinical staff and patients in various ways: making medical notes more understandable and directing patients' attention to relevant clinical results (Kambhamettu et al., 2024), recognizing out-of-hospital cardiac arrest during emergency calls (Blomberg et al., 2021), or predicting the progression of chronic kidney disease (Zheng et al., 2024). As

AI systems grow more capable, policymakers, healthcare providers, and developers stress the need for continuous ethical scrutiny. The European Commission's ethics guidelines for trustworthy AI note that individuals "interacting with AI systems must be able to keep full and effective self-determination over themselves" (AI HLEG, 2019, p. 12). Essentially, such complete self-determination reflects an idea of autonomy, i.e. an ability to "self-rule that is free from both controlling interference by others and from limitations, such as inadequate understanding, that prevent meaningful choice" (Beauchamp & Childress, 2001, p. 58).

Empirical studies have already explored the interaction between human autonomy and AI-assisted decision making. Here are three representative quotes:

> *"Although participants were open to the use of novel AI solutions in their care, most of them, especially those with diabetes, would prefer for them to be implemented alongside clinicians to keep human control and doctor-patient relationship in medicine"* (Schaarup et al., 2023, p. 6).
>
> *"These results highlight a potential concern that AI will impact the therapeutic alliance and potentially reduce the positive relational elements (i.e. empathy, understanding, etc.) present within the doctor-as-a-person elements of patient-centered care"* (Witkowski et al., 2024, p. 8).
>
> *"Many respondents at some point entertained the thought of the AI, at some point, being able to completely automate the tasks of the physicians, but always immediately observed that this would be a very bad evolution"* (Van Cauwenberge et al., 2022, p. 8).

These comments all illustrate a common ethical stance: AI should not exercise the kind of "controlling interference" (Beauchamp & Childress, 2001, p. 58) that would replace or undermine human physicians' final say in clinical decisions. Indeed, human autonomy is in these quotes seen as crucial for a successful integration of AI in healthcare, both to preserve important existing social relationships between doctor and patient and to avoid the undesirable replacement of human physicians, even if technically feasible.

We can think of such ethical stances as establishing broad, overarching ideals that are intended to apply universally across clinical practices while deliberately avoiding a case-by-case focus on individual circumstances (Beauchamp & Childress, 2001). This approach typically understands ethical principles *theoretically*: as principles that can be reasoned about abstractly rather than derived from empirical realities. As Bærøe et al. (2017, p. 7) argue, the aim of theoretical ethics is to develop "general ethical principles for medicine and health promotion" without direct reference to specific human circumstances. Principles of theoretical ethics impose non-negotiable requirements, as seen, for example, in AI governance frameworks (AI HLEG, 2019), where conditions for AI use are set independently of potential benefits that might arise in particular cases (Haward & Janvier, 2019).

However, the idea that AI systems should only be granted a subordinate role to human physicians' decision-making has been criticized for being overly abstract and difficult to implement in practice (Mittelstadt, 2019). In response to this, researchers have turned toward the "nuanced reality regarding how patients' values, such as autonomy, manifest in the everyday" (Maris et al., 2024, p. 11), aiming to develop healthcare AI that is ethically acceptable in practice. This practice-first focus asks us to examine how ethical considerations are prioritized and enacted across different clinical contexts and cultures and for differing normative ends (Pols, 2015). For example, patients in acute health crisis may prefer for clinicians to take responsibility, while patients in other settings may in some situations value assuming responsibility for their own treatment (Pols, 2015). Yet we still only know too little about what concretely makes AI ethically acceptable in particular clinical settings – knowledge that is essential if growing reliance on AI is to remain aligned with core ethical principles.

In this paper, we pursue a theoretical-philosophical inquiry (Bærøe & Bringedal, 2022) into the widely held view that human autonomy must always outrank healthcare AI for the latter to be ethically acceptable. We start with what we call the *Non-Autonomous AI Principle*. Roughly, according to this principle, healthcare AI used for decision-support should never possess the authority of human clinicians and should always be subject to meaningful human override. Because this principle – or, more carefully as we shall see, this bundle of ideas – is typically presented as applying across clinical contexts, we treat it as a principle of theoretical ethics. We argue, however, that there are concrete clinical practices in which AI systems are deployed and regarded as ethically acceptable in ways that put pressure on the Non-Autonomous AI Principle. In such settings, the principle's claim to universal authority is contested by how care is organized and justified. Our argument proceeds in three steps.

First, we survey research and policy concerns about autonomous AI systems in healthcare and relate them to familiar autonomy ideals in biomedical ethics, distilling these into the Non-Autonomous Principle (Section 2). Second, drawing on policy and press documents, we examine Danish breast cancer screening, where AI is increasingly deployed as first reader of mammograms (Section 3). This setting is relevant for testing the Non-Autonomous AI Principle because, in practice, the first reader AI has an authority comparable to a radiologist: its assessments carry equal weight and cannot be unilaterally overruled by the human second reader. Moreover, since this use of healthcare AI is by now common in Denmark, it hence also carries a seeming ethical legitimacy from its routine use in clinical practice and institutional promotion.

Third, we probe two practice-grounded explanations (Section 4) for why violating the Non-Autonomous AI Principle can nevertheless be ethically acceptable in breast cancer screening, namely *structural* factors such as shortages of radiologists, high throughput demands, as well as standardized workflows, and *existential* factors tied to the emotional and psychological states of cancer screening, diagnosis, and treatment. While structural challenges are common across many areas of healthcare, we argue that existential concerns may be a decisive factor in why autonomous first reader AI is ethically acceptable in breast cancer screening but not necessarily in other medical contexts. Taken together, these admittedly partial considerations specifically related to clinical practice (Pols, 2015) provide reasons to doubt that the Non-Autonomous AI Principle can serve as a universal standard for ethically acceptable AI use.

As such, the breast cancer screening case shows that the idea that healthcare AI must always remain strictly subordinate to, and available for overriding by, human clinicians, cannot serve as an ethical ideal across all contexts. Recognizing this allows us to develop ethically acceptable healthcare AI grounded in the actual intra-normativities of practice – how values are prioritized, negotiated, and enacted in specific workflows – rather than in abstract, one-size-fits-all ethical principles.

Before proceeding, two clarifications about the labels "theoretical ethics" and its opposite "empirical ethics" are in order. First, we do not claim that proponents of theoretical ethics must endorse the Non-Autonomous AI Principle, or that they must treat autonomy as a prior value that always overrides other moral considerations. A consequentialist may work entirely within a theoretical framework and nonetheless reject any strict priority for autonomy or any rule-or constraint-based moral theory. What we argue is that much existing AI ethics and governance discourse in healthcare presents requirements that protect human autonomy as if they were universal constraints on acceptable AI use. The Non-Autonomous AI Principle is meant to capture this influential strand of theory, not to exhaust the space of theoretical approaches.

Second, we do not assume that "empirical ethics" must oppose strict rules or principles. Empirically informed approaches can yield relatively rigid norms. The difference is that these norms are reconstructed from within practices rather than both stipulated prior to and defended largely independently of implementation and negotiation in concrete settings. Likewise, theoretical approaches can, and often do, take empirical realities seriously. But their starting point typically lies elsewhere: They begin with principles, rules, or ideals which are then applied to practice. In what follows, we use "theoretical" and "empirical" as ideal-typical starting points rather than as mutually exclusive camps. The former is principle-first, abstracting from practice; the latter is practice-first, grounded in intra-normativity. We treat the Non-Autonomous AI Principle as a prominent example of the former and our core claim is that an intra-normative, practice-first analysis of breast cancer screening explains why its assumed universality should be questioned.

The Non-Autonomous AI Principle as Foundational for Healthcare AI

To develop a precise statement of the Non-Autonomous AI Principle, let us begin by examining the concept of autonomy in healthcare and its relevance to AI. Suppose that, as a healthcare professional, you visit an elderly woman in deteriorating health. She has become delirious and refuses to leave her bed. You urge her to go to the hospital, but she refuses. According to modern medical ethics training, you should prioritize the elderly woman's autonomy. The woman should not be coerced into care against her will even if you believe hospitalization is in her best interest. Scenarios like this are commonly used to teach medical students the importance of autonomy (Hébert et al., 1992). In their landmark work on biomedical ethics, Beauchamp and Childress (2001) characterize respect of autonomy as:

> *"To respect an autonomous agent is, at a minimum, to acknowledge that person's right to hold views, to make choices, and to take actions based on personal values and beliefs. Such respect involves respectful action, not merely a respectful attitude. It also requires more than noninterference in others' personal affairs. It includes, at least in some contexts, obligations to build up or maintain others' capacities for autonomous choice while helping to allay fears and other conditions that destroy or disrupt their autonomous actions"* (Beauchamp & Childress, 2001, p. 63).

In the case of the elderly woman, we must thus first must recognize her right to refuse to go to the hospital. We must also ensure that the elderly woman has the capacity to make an informed decision, not impaired by cognitive ailments, rather than passively accepting any decision the woman makes.

This reasoning highlights why respect for autonomy serves as a paradigmatic example of a theoretical ethical principle: It is articulated as a general normative commitment designed to apply across clinical contexts regardless of cultural, institutional, or situational variation. It is not reconstructed from within particular practices but advanced as a general guiding principle for what clinicians ought to do – even where it may conflict with their own judgment or context-sensitive concerns.

Pellegrino (1994) argues that autonomy is equally critical for physicians. He describes the "physician-patient relationship [as, ed.] one of mutual obligation – like any truly ethical relationship" (Pellegrino, 1994, p. 51). Accordingly, if autonomy is central to ethical patient care, it also extends to physicians. Fostering their capacity for autonomous choice and recognizing their right to act on personal and professional judgment is part of the ethical principle of autonomy. A physician's autonomy is rooted in their individuality as a person, their expertise as a medical professional, and their role within a broader ethical and professional community (Pellegrino, 1994). Thus, autonomy in biomedical ethics also concerns the autonomy of physicians.

Respect for Autonomy

Given the dominant role that autonomy plays in biomedical ethics, it is natural to hold that healthcare AI should also respect the autonomy of both patients and physicians. Roughly, this means that AI should support their capacity for informed and voluntary judgment based on reason, enabling both patients and clinicians to form their own assessments, act on what they take to be right, and retain meaningful control over decisions.

Unlike patients and physicians, however, while AI may "embody certain values, [...] it does not embody them to the extent that it can be held responsible for them" (Sheir et al., 2024, p. 12). On the standard view in biomedical ethics, autonomy is closely linked to responsibility: We treat a patient or physician as autonomous not merely because their behavior reflects certain values, but because they can be asked to justify choices, answer for consequences, and revise their stance. To be autonomous is, at least partly, to be an appropriate target of demands for justification and practices of praise or blame. Given that AI systems neither hold values nor make choices in a way that grounds responsibility, they cannot possess autonomy in the same sense as humans. In this sense, while AI should respect patient and physician autonomy, we are not seemingly ethically obligated to respect "AI autonomy" itself.

This ethical asymmetry is also evident in clinical practice. Consider IBM's Watson Oncology, an AI system that offered treatment recommendations for cancer patients. This system was piloted in multiple countries, including Denmark, where it was quickly found to issue potentially dangerous advice (Djursing, 2017). As a result, the trial was quickly stopped. Crucially, no one asked the system underlying Watson Oncology to justify its "values" or "beliefs". Had a physician made comparable errors, they would almost certainly have been required to account for their decisions, possibly retrained, and held responsible for the consequences. This contrast underscores the point above: While we expect AI to support and respect human autonomy, we do not treat AI itself as an autonomous moral or professional agent. It is not held to practices of justification and it is not a target of blame or praise. When an AI system fails, we simply turn it off (and investigate the humans and institutions around it).

Because AI lacks the kind of responsibility assumed in clinicians' autonomy, its proper role is hence supportive and under human oversight. This claim is echoed in broader healthcare discussions. For example, in a bulletin from the American College of Radiology, chief medical officer Nina E. Kottler cites radiology professor Curt Langlotz as predicting that "radiologists are not going to be replaced by AI, but radiologists who use AI will replace those who do not" (Hudnall, 2021). Policy bodies likewise emphasize sustained human oversight (AI HLEG, 2019; European Commission, 2021). Empirically, a systematic review in oncology finds that most FDA-cleared AI devices function as tools for supporting clinicians rather than autonomous decision-makers (Luchini et al., 2022). Together, these trends follow naturally from the autonomy gap: AI is introduced to assist human judgment, not to substitute it.

Human Oversight

So why should healthcare AI remain subject to human oversight? Drawing on recent research, there are at least three key reasons.

The first reason relates to errors and bias. Because AI systems learn statistical patterns from training data, they will inevitably make errors in their classifications – even in settings that look routine. This risk is amplified when no physician reviews the system's outputs (Filippi et al., 2023; Najjar, 2023). Unlike physicians, AI does not reflect on and correct its own mistakes during its decision-making which can cause frustration among healthcare personnel (Beede et al., 2020). Evidence also indicates that physicians are more forgiving of human errors than AI errors (Lenskjold et al., 2023). At the same time, reliance on AI may introduce automation bias, nudging clinicians to over-trust AI-generated outputs (Bernstein et al., 2023; Fang et al., 2024). In AI-assisted breast cancer screening, this may for instance lead to unnecessary and costly patient recalls (Hernström et al., 2025).

Second, there are reasons to do with maintaining critical feedback loops and professional development for clinicians. For instance, as a second reader in breast cancer screening, AI systems can surface alternative hypotheses, counter physician bias, and prompt review and critique. Such benefits, which arguably constitute an important factor in the ethical acceptability of AI systems, would be significantly weakened if they could operate autonomously without human oversight. Empirically, Bergquist et al. (2023) report that radiologists were more likely to trust reliable AI when it serves as a second opinion that helps reduce radiologists' bias, especially for junior physicians. Clinicians also envision using AI to cross-check and refine their own judgments (Fang et al., 2024) and to demonstrate expertise to colleagues (Beede et al., 2020). Even in cases where AI matches human performance, fully replacing physicians with AI hence risks removing essential training opportunities for junior doctors (Kühl et al., 2023).

Third, there are reasons concerning safeguarding the therapeutic alliance between patients and physicians. When AI systems shift from assisting to replacing elements of clinical practice, patients can feel less understood and more distant from their caregivers (Mack et al., 2009). Empirical work echoes this concern: In a workshop with patients, Adams et al. (2020) found that participants worried radiology AI would lead to reduced face-to-face interaction with physicians. While many may be comfortable with AI as a supportive tool or source of recommendation, they worry about losing human connection (Witkowski et al., 2024) and being treated as data points rather than whole persons (Maris et al., 2024).

No Final Authority

The preceding arguments already speak in favor of keeping AI systems in a purely supportive role. But there is a further independent consideration that is worth mentioning, namely that, in clinical care, AI should not be granted final authority. First, on grounds of accountability, someone must be able to answer for AI-generated errors. To avoid "responsibility gaps", scholars caution against autonomous AI and argue for retaining a human decision-maker (Najjar, 2023; Sheir et al., 2024) . In radiology, for instance, Mudgal and Das (2024, p. 7) contend that establishing "legal liability and responsibility" requires keeping radiologists in the loop so AI-propagated mistakes do not create dead ends where no party can be held to account.

Second, on grounds of explainability, many AI systems are difficult – if not impossible – to fully explain (Najjar, 2023). Patients themselves stress this. In interviews with individuals at risk of sudden cardiac death, participants emphasized the need for clinicians to understand and judge the quality of AI outputs (Maris et al., 2024). Likewise, Lenskjold et al. (2023) argue that clinicians must be able to interrogate AI's reasoning to spot and mitigate possible errors. In short, precisely because AI can err and often resists full explanation, clinicians should remain the final responsible parties for administering care.

The Non-Autonomous AI Principle

So, there are significant ethical and practical reasons to be skeptical about autonomous AI systems in healthcare, reinforcing the need to limit their decision-making capabilities. Drawing on the literature above, we can reconstruct a cluster of related commitments often treated as a single package. For our purposes, it is helpful to distinguish three components:

1. *Autonomy-Respect Condition*: AI systems should respect human autonomy. They should not unjustifiably undermine patients' or physicians' ability to form their own informed and voluntary judgments or act on their values.
2. *Human Oversight Condition*: AI systems should be embedded in workflows that preserve meaningful human oversight. Clinicians must be able to monitor, interrogate, and challenge AI outputs rather than passively accepting AI recommendations.
3. *No-Final-Authority Condition*: AI systems should not occupy roles of final, non-reviewable authority in clinical decision-making. They should not be granted the same formal decision-making standing as human clinicians nor function as sole decisive gatekeepers for diagnosis or treatment.

For ease of presentation, we refer to this three-part package as the *Non-Autonomous AI Principle*, though "principle" may not capture how these components are analytically separable. One could, in principle, endorse the Autonomy-Respect Condition while rejecting the No-Final-Authority Condition, or insist on oversight while allowing some limited forms of AI decision-making authority. We do not claim that respect for autonomy entails either human oversight or rejection of any autonomous AI. Rather, as motivated above, our claim is that much existing governance and ethics discourse on healthcare AI treats these three ideas collectively, presenting them as if they formed a single, unified constraint on acceptable AI use.

Consider again the delirious elderly woman who refuses hospital admission. Suppose an AI system, using healthcare data such as blood tests, recommends admission for further evaluation. The Autonomy-Respect Condition requires that the patient's will and capacity for autonomous choice remain central, and the practitioner must assess whether her refusal reflects a competent, non-compromised decision. The Human Oversight Condition demands that the physician critically examines the AI systems' recommendation considering her clinical judgment and the patient's situation, rather than treating AI output as settled. Finally, the No-Final-Authority Condition rules out granting the AI sole power over deciding the patient's admission, which ensures that the decision rests with the clinician.

Thus, the cluster of commitments in the Non-Autonomous AI Principle preserves the ethical centrality of human patients and clinicians in standard healthcare settings. Crucially, this seems to hold *even when* AI systems outperform humans on accuracy. A randomized clinical trial testing ChatGPT's effect on physicians' diagnostic reasoning found that ChatGPT alone outperformed both physicians working independently and physicians assisted by ChatGPT (Goh et al., 2024). Yet, the study's authors explicitly caution that this does not suggest that "LLMs should be used for diagnosis autonomously without physician oversight" (Goh et al., 2024, p. 7). Similarly, while AI systems can be effective as first reader in radiology, replacing human radiologists in that role would likely widen skill gaps and generate long-term professional challenges (Kühl et al., 2023). In short, unlike the case of IBM's Watson Oncology where autonomy was reduced because of evidence errors, AI autonomy remains ethically unacceptable even when the system demonstrates superior accuracy. Oversight, accountability, and the clinician's decisional role still do normative work that accuracy alone cannot replace.

Accordingly, the Non-Autonomous AI Principle can be read as the paradigmatic response to autonomy concerns in both the healthcare-AI literature and policy. It states a default constraint that is often presumed to hold across clinical domains. While we do not deny the intuitive force of these commitments, our question is how far they can reasonably be regarded as universal. As we shall see in the next section, once we take the intra-normative features of particular practices seriously, the claim to universality becomes doubtful.

Breast Cancer Screening with AI and the Intra-normativity of Autonomy

We now examine a case of AI-driven breast cancer screening that we argue conflicts with the plausibility of the Non-Autonomous AI Principle. We begin by describing the case, then outline its relation to the Non-Autonomous AI Principle, and finally explore how the case challenges the Non-Autonomous AI Principle.

The Case – AI as First Reader in Breast Cancer Screening

In Denmark, breast cancer screening is a politically governed screening process, with the Danish Health Authority setting screening guidelines and recommendations. As part of a national strategy to reduce breast cancer mortality, women living in Denmark aged 50-69 are invited to undergo breast cancer screening regardless of symptoms. The screening process involves performing X-rays (mammograms) of both breasts (Danish Health Authority, n.d.). This procedure takes approximately five minutes, and women do not have the opportunity for further consultation during the examination (Region Midtjylland, n.d.).

After imaging, mammograms are analyzed in a two-step reading process. Two radiologists, a first and a second reader, independently assess the scans to determine whether abnormalities suggestive of breast cancer are present. This blinded review ensures that neither reader is influenced by the other's assessment. If the two radiologists disagree, a consensus conference is held where a third reading is conducted to determine whether further diagnostic procedures are needed (Elhakim et al., 2020). If additional testing is required, the woman is invited for further examinations, such as additional imaging, to assess the actual presence of breast cancer (Region Midtjylland, n.d.). Note that the goal at this stage is not to diagnose breast cancer but rather to find any abnormalities that might signal the possibility of breast cancer (Region Midtjylland, n.d.).

In response to a shortage of radiologists in Danish cancer care (The Boston Consulting Group, 2017) and the time-intensive nature of breast cancer screening and diagnostics (Kommission for robusthed i sundhedsvæsenet, n.d.), the Capital Region of Denmark tested an AI system to function as a first reader in breast cancer screening (Region Hovedstaden, 2023). The system analyzes mammograms and assigns a risk score (0–100), classifying cases as low risk (0–42), medium risk (43–74), or high risk (75+). This score is displayed in a user interface alongside a descriptive label – for example, a risk score of 95 is labeled as "Elevated" (Kommission for robusthed i sundhedsvæsenet, n.d., p. 372). In this workflow, the

AI system replaces human first readers for mammograms classified as low risk (Danske Regioner, 2024). The Regions of Central and Southern Denmark have since launched a joint initiative to "investigate how you can most safely implement the use of AI in the breast cancer screening program" (Region Midtjylland, 2023, p. 4). The role of AI in breast cancer screening was further politically solidified in June 2024 when a report recommended expanding the use of AI in healthcare based on its success in breast cancer screening (Sundhedsstrukturkommissionen, 2024).

Since employing AI-driven systems as a first reader AI in breast cancer screening is hence politically accepted as a way to enable "good patient care" (Pols, 2015, p. 87), these systems are inside the domain of ethically acceptable practices. But what justifies their ethical acceptability, and does this justification align with the Non-Autonomous AI Principle?

An AI First Reader and the Non-Autonomous AI Principle

Let us first consider how the Non-Autonomous AI Principle applies to the case of first reader AI in breast cancer screening. Despite their naming, first and second readers do not assess mammograms sequentially. Rather, they conduct separate, independent evaluations without knowledge of each other's assessments. So, imagine an AI system that replaces the first reader and achieves a level of diagnostic accuracy that is comparable to human radiologists (Hernström et al., 2025; Kühl et al., 2023). Because the human second reader is blinded to the first reader's assessment, they have no way of knowing whether the first reader is a human or an AI system, nor what the first reader's assessment was. Furthermore, both assessments carry equal weight when it comes to deciding whether a third reading is needed. That is, if either reader detects abnormalities while the other does not, the case moves to a consensus conference regardless of whether the first reader was AI or human (Bugge, 2023).

This practice, we suggest, challenges each element of the Non-Autonomous AI Principle. For the first condition concerning respect for autonomy, note that the radiologists' professional judgment at the first reading stage is structurally constrained: they are required to act on disagreement with the AI first reader. That is, when it comes to radiologist autonomy, the second reader is structurally required to treat the AI's assessment as an independent judgment containing recognized institutional consequences, i.e. as triggering a consensus conference upon disagreement. This workflow thus prevents the radiologist from simply treating the AI system as an imperfect tool that can be ignored in a specific case. Professional discretion gets limited by procedure rather than by medical judgment alone. Thus, the physicians' ability to act on their values and professional judgment is partly displaced by a non-answerable, non-responsible system and the rules built around it. Likewise, when it comes to patient autonomy, patients encounter screening as a highly standardized procedure with little room for expressing individual values. Since the introduction of the relevant AI system is explicitly justified in terms of efficiency and speed, its inclusion does not speak to any richer forms of patient self-determination.

With respect to the second condition, meaningful human oversight of the AI's assessments is arguably weakened. Because the reading is blinded, the human second reader does not see the AI's assessment. This means that they cannot easily monitor, interrogate, or challenge that assessment on a case-by-case basis. Effectively, they cannot question the AI system's decision. Human oversight is hence only exercised indirectly and after the event of reading at the overall level of workflow design – for instance by institutionalizing that disagreement between first and second reader triggers a consensus conference. But where individual AI recommendations are made and acted upon, clinicians are reduced to following general rules and principes of the workflow design. And this, we claim, comes very close to constituting non-meaningful oversight in any typical sense.

For the third condition, the AI reader's role in the screening process also sits uneasily with the idea that AI systems should not possess authority to make decisions within clinical workflows. During first reading, the AI's judgments have a fixed institutional weight equivalent to a human reader's. It functions as an independent "vote" that cannot simply be set aside, and disagreement with a human radiologist automatically triggers the move to consensus conference. In low-risk cases, the AI fully replaces the human first reader, determining which women are routed onward in the screening pathway and which are not. While human clinicians retain ultimate responsibility at later stages of care, the AI system exhibits a locally authoritative gatekeeping role at a crucial decision point. This reflects the kind of role that the third element of the Non-Autonomous AI Principle is meant to exclude.

Accordingly, given the use of AI in breast cancer screening is regarded as ethically acceptable – which current practice clearly supports – these observations cast doubt on the Non-Autonomous AI Principle's universal applicability. Concretely, autonomy is reshaped by binding both patients and physicians into highly standardized AI-driven routines; meaningful human oversight is minimized at the level of individual cases; and AI's application in the workflow functions as a decision-maker and gatekeeper. Of course, not every technical substitution for a human function diminishes autonomy or responsibility. In many cases, like the use of routine measuring instruments, functional substitution leaves the underlying normative structure largely intact. What the breast cancer screening case shows, however, is that *this particular implementation* of AI places significant pressure on all three elements of the Non-Autonomous AI Principle – above all, the prohibition on AI occupying roles with decisional power comparable to that of human.

Thus, the breast cancer screening practice threatens the Non-Autonomous AI Principle as a universal constraint. If all three of its core elements can be relaxed in this setting without undermining the perceived ethical legitimacy of the practice, then the principle cannot straightforwardly serve as an overarching standard for evaluating healthcare AI. In clinical practice, AI first readers are integrated into decision-making structures as locally authoritative decision-makers like human readers, shaping how human autonomy and responsibility are exercised in breast cancer screening.

The case also exposes a deeper mismatch between a principle-first, theoretical approach to ethics and an "empirical" ethics grounded in intra-normativity and empirical practice. Much theoretical work articulates the commitments that we collect in the Non-Autonomous AI Principle by drawing a relatively strict normative boundary between human and machine agency – presuming that AI should remain a subordinate tool under human control. By contrast, from an intra-normative view (Pols, 2015), ethics is reconstructed from within clinical routines and infrastructures. Here, practice reveals that AI is *already* embedded in workflows with capacities comparable to human clinicians: AI first readers do not merely advise but issue assessments that steer diagnostic pathways, trigger consensus conferences, and cannot be unilaterally overridden.

In breast cancer screening, then, AI first readers count not simply as departures from a theoretical rule. They reflect how autonomy is negotiated and enacted within a specific clinical setting. To the extent that integrating AI first readers enables good care, intra-normative processes thus help shape new forms of ethical legitimacy for AI. Here, AI systems are no longer mere instruments in the medical toolbox but integral parts of decision-making structures, blurring boundaries between human and machine agencies while reshaping how responsibility and autonomy are exercised in practice.

The Intra-normativity of Autonomy

It seems that deviations from the Non-Autonomous AI Principle are not only tolerated but also ethically acceptable. How can we explain the source of this?

Structural Issues

One way to approach this question is by examining the structural challenges that have shaped the adoption of AI in breast cancer screening. The testing of first reader AI in the Capital Region of Denmark has received widespread attention in both press and governmental reports where breast cancer screening is consistently framed as struggling under a "lack of employees, among these especially those consulting physicians who review and assess" mammograms (Bugge, 2023). This shortage is exacerbated by the fact that mammogram analysis is often performed as an additional task with radiologists frequently using spare time to keep up with screening demands. Compounding this problem is an increase in referrals to breast cancer screening among asymptomatic Danish women (Region Midtjylland, 2023), which creates bottlenecks for patients who might actually exhibit symptoms of breast cancer (Region Hovedstaden, 2023).

Breast cancer has received widespread attention as an arena for AI precisely due to these structural issues. In the North Jutland Region's press release on AI in breast cancer screening (Hyldgaard, 2024), the Region highlights persistent bottlenecks – not least a shortage of radiologists with breast-imaging expertise. As above, director of a radiological unit Peter Buss Lasborg notes that "something we are looking at in the future is artificial intelligence in relation to prostate cancer" (Hyldgaard, 2024), envisioning systems that can draw tumor locations on radiological images. The takeaway is that breast cancer screening is an obvious priority for AI adoption in healthcare, being valued partly for its potential to ease staff shortages by reducing waiting times and relieving time pressure. But that promise depends on performance at the point of care. First reader AI must be both fast and reliable; if it is slow or faulty, it will fail to address those structural problems and may even disrupt clinical workflows.

By framing AI as a solution to an overstretched system, the ethical acceptability of first reader AI becomes intertwined with the practical need to sustain screening programs. But structural challenges alone cannot fully explain the ethical legitimacy granted to first reader AI. After all, the whole Danish healthcare system faces mounting pressures from demographic shifts, health inequalities, and economic constraints (Jenvall, 2024; Sundhedsstrukturkommissionen, 2024).

If structural challenges alone justified using autonomous healthcare AI, we would expect similar systems to be implemented across the entire Danish healthcare system, which has not occurred. This suggests that something beyond efficiency concerns shapes the ethical acceptance of AI autonomy in breast cancer screening.

Existential Concerns

Another key factor, we suggest, concerns the existential distress associated with waiting for a potential cancer diagnosis. Several stakeholders in Danish cancer care highlight the fear and anxiety that patients often experience while waiting for a diagnosis. For example, the Capital Region of Denmark, discussing the pros and cons of breast cancer screening, identifies the following concern:

> *"False alarm: If the X-ray images show changes that the physicians assess can be signs of cancer, you will be invited to a new examination. In some instances, it turns out to be the case that the changes are benign, meaning it was a false alarm"* (Region Hovedstaden, n.d.).

This quote highlights how breast cancer screening can lead to distressing "false alarms" where women become worried about a possible cancer diagnosis later identified as unfounded. Screening does not definitively confirm cancer but indicates when further testing is warranted. This inherent uncertainty entails considerable emotional strain (Danish Health Authority, 2023). The Danish Cancer Society's guide for patients coping with anxiety notes that the mere mention of cancer can provoke intense worry, and that hospital-based examinations like screening may themselves cause distress (Aglund, 2024). In this context, speed is thus tied directly to patients' emotional well-being.

Offersen et al. (2018) describe "cancer mythologies": narratives that shape how people make sense of cancer in their lived experience. One dominant mythology is cancer's "ghostly presence" where the disease is imagined as invisible, hidden, and silently growing in the body (Offersen et al., 2018, p. 39). Denmark's strong emphasis on early detection, even before symptoms appear, reinforces this view of cancer as a potentially ever-present threat. Therefore, patients may come to scrutinize their own bodies with a sense of unease and suspicion (Offersen et al., 2018). Robb et al. (2014) similarly found participants instinctively associated cancer with intense negative emotions, with one associating it with "[f]ear and confusion… generally it ends in death, catastrophe" (Robb et al., 2014, p. 2).

These studies suggest that moving through cancer care testing, diagnosis, treatment, and follow-up triggers emotional turmoil. Frumer et al. (2021) describe how patients may experience this process as a kind of suspended state, "being in the meantime" (Frumer et al., 2021, p. 19). This meantime is filled with ambiguities shaping how both patients and their relatives structure their lives. Uncertainty about a potential diagnosis can trigger both fear and apathy in patients and relatives (Frumer et al., 2021). Moreover, receiving a cancer diagnosis is an unsettling experience, with patients reporting how "[cancer, ed.] and the prospect of extensive treatment evoked a profound fear, leading to existential distress" (Wilhøft Kristensen et al., 2024, p. 3).

In this light, autonomous first reader AI for breast cancer screening may be desirable precisely because it can minimize the existential fear and anxiety associated with waiting. Faster screenings mean less time spent in the meantime, relieving distress for women awaiting results and allowing for more efficient allocation of resources to those requiring follow-up examinations. The structural challenges discussed earlier, including workforce shortages, increasing patient referrals, and bottlenecks in screening programs, exacerbate the meantime by prolonging waiting periods for both routine screenings and diagnostic follow-ups. By accelerating the process, autonomous first reader AI directly addresses these challenges, reducing delays and lightening patient distress. Thus, it becomes an ethically desirable form of care mitigating both practical and existential burdens in breast cancer screening.

Concluding Remarks

We began this paper with a call to ground healthcare AI ethics in practical, context-sensitive considerations. We then reconstructed the Non-Autonomous AI Principle as a cluster of widely held commitments in theoretical ethics and AI governance: Healthcare AI should respect autonomy, remain under meaningful human oversight, and not occupy roles with decisional standing comparable to clinicians. Our analysis of first reader AI in the breast cancer screening workflow suggests that, for this case, each element is strained. Autonomy is reshaped, oversight is thinned out at the level of individual cases, and AI acquires a locally authoritative gatekeeping role. Despite this, the practice is not understood as ethically unacceptable from within. This indicates that the Non-Autonomous AI Principle – at least in its strongest and most universal form – cannot straightforwardly function as a general ethical standard for evaluating healthcare AI.

Considered through empirically oriented ethics, the implications become clear. Rather than relying solely on general, principle-first accounts, we also need a framework that takes seriously the norms and values of specific clinical practices. Through an intra-normative lens, a functionally autonomous first reader AI in breast cancer screening is not only acceptable but may even constitute a normative good – given structural pressures on screening programs and the existential burdens of waiting for a possible cancer diagnosis. Importantly, this insight does not automatically generalize across medicine. As Pols (2015, p. 87) notes, ethical acceptability takes form only when we "compare values and contexts". Our case should therefore be read as an invitation to examine, domain by domain, how the elements of the Non-Autonomous AI Principle are potentially reconfigured – sometimes upheld, sometimes relaxed – across different areas of healthcare.

To illustrate this sensitivity to context, consider the contrast between breast cancer screening and sudden cardiac death (SCD) prevention. Maris and colleagues (2024) show that in SCD prevention, patients place high value on human physicians' presence – especially for assuming responsibility for medical decisions and cultivating a therapeutic alliance – and they recognize the subjective elements of medical judgment (e.g., whether to implant a pacemaker). In Pols' (2015) terms, good care here involves clinicians taking responsibility onto themselves and walking patients through uncertainty. Ethically acceptable AI for SCD prevention must therefore respect this form of autonomy and responsibility.

By contrast, breast cancer screening is organized very differently. It is fast, standardized, and largely without direct physician-patient interaction. Typical interactions last minutes, with no conversation with a clinician. The care need is hence different from SCD prevention; it involves informing and reassuring patients without amplifying existential distress to "foster a feeling of safety" (Pols, 2015, p. 87). In this setting, as we have seen, an autonomous first reader AI can be ethically acceptable – indeed potentially beneficial – provided it supports timely, consistent triage and reduces unnecessary recalls while maintaining safety.

Intra-normativity can help explain why autonomy and oversight are configured differently across practices. What is ethically appropriate for SCD prevention may not transfer directly to breast cancer screening. The same practice-first perspective may clarify other medical domains. For instance, Mainz et al. (2023) argue that in AI-assisted sepsis detection, speed can ethically outweigh modest accuracy gains because harms of delayed treatment may exceed the relatively low costs of overtreatment (e.g., unnecessary distributions of antibiotics). Like before, "good care" may prioritize crisis-prevention, making fast AI ethically preferable, much as in our screening case.

By aligning AI systems with the concrete norms, care need, and workflow pressures of specific clinical contexts, we can thus design ethically acceptable AI systems without leaning on abstract one-size-fits-all principles. Still, autonomous healthcare AI raises difficult questions about oversight, accountability, and the kinds of errors systems are prone to (Bjerring et al., 2025). Future work should map where autonomy supports good care and where it threatens it, clarify role-responsibilities across clinicians, institutions, and developers, and create validation regimes that track clinical value and not just technical metrics. Only then can AI-assisted healthcare be both ethically sound and practically viable.

References

Adams, S. J., Tang, R., & Babyn, P. (2020). Patient Perspectives and Priorities Regarding Artificial Intelligence in Radiology: Opportunities for Patient-Centered Radiology. *Journal of the American College of Radiology*, 17(8), 1034–1036. https://doi.org/10.1016/j.jacr.2020.01.007

Aglund, M. (2024, July 6). *Angst ved en kræftsygdom—Kræftens Bekæmpelse*. Kræftens Bekæmpelse. https://www.cancer.dk/hverdag/dit-liv-med-kraeft/reaktioner-og-tanker/angst/

AI HLEG. (2019). *ETHICS GUIDELINES FOR TRUSTWORTHY AI.* European Commission. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

Bærøe, K., & Bringedal, B. (2022). Social inequalities in health: The need for a new concept of health. *Tidsskrift for Forskning i Sygdom Og Samfund*, 19(36), 27–38. https://doi.org/10.7146/tfss.v19i36.133048

Bærøe, K., Ives, J., De Vries, M., & Schildmann, J. (2017). On classifying the field of medical ethics. *BMC Medical Ethics*, 18(1), 30, s12910-017-0193–x. https://doi.org/10.1186/s12910-017-0193-x

Beauchamp, T. L., & Childress, J. F. (2001). P*rinciples of Biomedical Ethics* (5th ed.). Oxford University Press. https://ebookcentral-proquest-com.ez.statsbiblioteket.dk/lib/asb/detail.action?docID=5763592

Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., & Vardoulakis, L. M. (2020). A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12. https://doi.org/10.1145/3313831.3376718

Bergquist, M., Rolandsson, B., Gryska, E., Laesser, M., Hoefling, N., Heckemann, R., Schneiderman, J. F., & Björkman-Burtscher, I. M. (2023). Trust and stakeholder perspectives on the implementation of AI tools in clinical radiology. E*uropean Radiology*, 34(1), 338–347. https://doi.org/10.1007/s00330-023-09967-5

Bernstein, M. H., Atalay, M. K., Dibble, E. H., Maxwell, A. W. P., Karam, A. R., Agarwal, S., Ward, R. C., Healey, T. T., & Baird, G. L. (2023). Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. *European Radiology*, 33(11), 8263–8269. https://doi.org/10.1007/s00330-023-09747-1

Bjerring, J. C., Busch, J., & Aastrup Munch, L. (2025). A Counterfactual Account of Algorithmic Robustness. *Minds and Machines*, 35(3), 34. https://doi.org/10.1007/s11023-025-09734-z

Blomberg, S. N., Christensen, H. C., Lippert, F., Ersbøll, A. K., Torp-Petersen, C., Sayre, M. R., Kudenchuk, P. J., & Folke, F. (2021). Effect of Machine Learning on Dispatcher Recognition of Out-of-Hospital Cardiac Arrest During Calls to Emergency Medical Services: A Randomized Clinical Trial. *JAMA Network Open*, 4(1), e2032320. https://doi.org/10.1001/jamanetworkopen.2020.32320

Bugge, M. (2023, July 16). *Hovedstaden får kunstig intelligens til at opdage brystkræft—Resten af landet vil gøre den kunsten efter.* DR. https://www.dr.dk/nyheder/viden/teknologi/hovedstaden-faar-kunstig-intelligens-til-opdage-brystkraeft-resten-af-landet

Danish Health Authority. (n.d.). *Information about the Danish breast cancer screening programme (Health for All)*. Danish Health Authority. Retrieved September 15, 2024, from https://www.sst.dk/-/media/Udgivelser/2023/Kraeft/Screening/UK-Pjece-brystkraeft.ashx

Danish Health Authority. (2023, October 10). *Screening for brystkræft.* Sundhedsstyrelsen. http://www.sst.dk/da/Borger/Sygdomme-og-lidelser/Kræftsygdom/Screening-for-kraeft/Brystkraeft

Danske Regioner. (2024). *Vi skal skabe et godt og effektivt sundhedsvæsen ved brug af kunstig intelligens.* Danske Regioner. https://www.regioner.dk/media/0urf11dv/kunstig-intelligens-udspil-070624.pdf

Djursing, T. (2017, October 23). *Doktor Watson i modvind: Foreslog livsfarlig medicin til danske patienter | Ingeniøren*. Ingeniøren. https://ing.dk/artikel/doktor-watson-i-modvind-foreslog-livsfarlig-medicin-til-danske-patienter

Elhakim, M. T., Graumann, O., Larsen, L. B., Nielsen, M., & Rasmussen, B. S. (2020). Kunstig intelligens til cancerdiagnostik i brystkræftscreening. *Ugeskr Læger*. https://ugeskriftet.dk/videnskab/kunstig-intelligens-til-cancerdiagnostik-i-brystkraeftscreening

European Commission. (2021). *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS.* EUR-Lex. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

Fajstrup, M. (2021, July 26). *AI og visuel feedback ruster lægerne til vanskelig hudkræftdiagnose | Version2*. https://www.version2.dk/artikel/ai-og-visuel-feedback-ruster-laegerne-til-vanskelig-hudkraeftdiagnose

Fang, N., Pluyter, J., Bakker, S., Jacobs, I., Luyer, M., Nederend, J., Raijmakers, J., Chen, L.-L., & Funk, M. (2024). From Experience to Experience: Key Insights for Improved Interaction with AI in Radiology. E*xtended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 1–7. https://doi.org/10.1145/3613905.3650779

Filippi, C. G., Stein, J. M., Wang, Z., Bakas, S., Liu, Y., Chang, P. D., Lui, Y., Hess, C., Barboriak, D. P., Flanders, A. E., Wintermark, M., Zaharchuk, G., & Wu, O. (2023). Ethical Considerations and Fairness in the Use of Artificial Intelligence for Neuroradiology. *American Journal of Neuroradiology*, 44(11), 1242–1248. https://doi.org/10.3174/ajnr.A7963

Frumer, M., Andersen, R. S., Vedsted, P., & Offersen, S. M. H. (2021). 'In the Meantime': Ordinary Life in Continuous Medical Testing for Lung Cancer. *Medicine Anthropology Theory*, 8(2), 1–26. https://doi.org/10.17157/mat.8.2.5085

Goh, E., Gallo, R., Hom, J., Strong, E., Weng, Y., Kerman, H., Cool, J. A., Kanjee, Z., Parsons, A. S., Ahuja, N., Horvitz, E., Yang, D., Milstein, A., Olson, A. P. J., Rodman, A., & Chen, J. H. (2024). Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial. *JAMA Network Open*, 7(10), e2440969. https://doi.org/10.1001/jamanetworkopen.2024.40969

Haward, M. F., & Janvier, A. (2019). Empirical Over Theoretical Ethics: Choosing What Matters to Patients and Families. *The American Journal of Bioethics*, 19(3), 54–56. https://doi.org/10.1080/15265161.2019.1572813

Hernström, V., Josefsson, V., Sartor, H., Schmidt, D., Larsson, A.-M., Hofvind, S., Andersson, I., Rosso, A., Hagberg, O., & Lång, K. (2025). Screening performance and characteristics of breast cancer detected in the Mammography Screening with Artificial Intelligence trial (MASAI): A randomised, controlled, parallel-group, non-inferiority, single-blinded, screening accuracy study. *The Lancet Digital Health*, S258975002400267X. https://doi.org/10.1016/S2589-7500(24)00267-X

Hudnall, C. (2021, June 1). *Thinking About AI?* American College of Radiology. https://www.acr.org/Practice-Management-Quality-Informatics/ACR-Bulletin/Articles/June-2021/Thinking-About-AI

Hyldgaard, S. (2024, April 2). *Kunstig intelligens opsporer brystkræft og afhjælper ressourcemangel i Region Nordjylland.* https://rn.dk/Service/Nyhedsbase-RN/Nyhed?id={625E1BB4-3512-410C-8023-B7E5136518E3}

Jenvall, L. (2024, June 1). *Ny økonomiaftale giver store smil, men regionerne er fanget i et krydspres.* DR. https://www.dr.dk/nyheder/politik/ny-oekonomiaftale-giver-store-smil-men-regionerne-er-fanget-i-et-krydspres

Kambhamettu, H., Metaxa, D., Johnson, K., & Head, A. (2024). Explainable Notes: Examining How to Unlock Meaning in Medical Notes with Interactivity and Artificial Intelligence. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–19. https://doi.org/10.1145/3613904.3642573

Kommission for robusthed i sundhedsvæsenet. (n.d.). *Robusthedskommissionens anbefalinger—Bilagsrapport.* Retrieved December 18, 2024, from https://www.ism.dk/Media/638336462433905130/Robusthed-Bilagsrapport-TILG.pdf

Kühl, J., Elhakim, M. T., Stougaard, S. W., Rasmussen, B. S. B., Nielsen, M., Gerke, O., Larsen, L. B., & Graumann, O. (2023). Population-wide evaluation of artificial intelligence and radiologist assessment of screening mammograms. *European Radiology*, 34, 3935–3946. https://doi.org/10.1007/s00330-023-10423-7

Lenskjold, A., Nybing, J. U., Trampedach, C., Galsgaard, A., Brejnebøl, M. W., Raaschou, H., Rose, M. H., & Boesen, M. (2023). Should artificial intelligence have lower acceptable error rates than humans? *BJR|Open*, 5(1), 20220053. https://doi.org/10.1259/bjro.20220053

Luchini, C., Pea, A., & Scarpa, A. (2022). Artificial intelligence in oncology: Current applications and future perspectives. *British Journal of Cancer*, 126(1), 4–9. https://doi.org/10.1038/s41416-021-01633-1

Mack, J. W., Block, S. D., Nilsson, M., Wright, A., Trice, E., Friedlander, R., Paulk, E., & Prigerson, H. G. (2009). Measuring therapeutic alliance between oncologists and patients with advanced cancer: The Human Connection Scale. *Cancer*, 115(14), 3302–3311. https://doi.org/10.1002/cncr.24360

Mainz, J., Munch, L., Bjerring, J. C., & Godtfredsen, S. (2023). Why algorithmic speed can be more important than algorithmic accuracy. *Clinical Ethics*, 18(2), 161–164. https://doi.org/10.1177/14777509221138750

Maris, M. T., Koçar, A., Willems, D. L., Pols, J., Tan, H. L., Lindinger, G. L., & Bak, M. A. R. (2024). Ethical use of artificial intelligence to prevent sudden cardiac death: An interview study of patient perspectives. *BMC Medical Ethics*, 25(1), 42. https://doi.org/10.1186/s12910-024-01042-y

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. https://doi.org/10.1038/s42256-019-0114-4

Mudgal, K. S., & Das, N. (2024). The ethical adoption of artificial intelligence in radiology. *BJR|Open*, 2(1), 20190020. https://doi.org/10.1259/bjro.20190020

Najjar, R. (2023). Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging. *Diagnostics,* 13(17), 2760. https://doi.org/10.3390/diagnostics13172760

Offersen, S. M. H., Risør, M. B., Vedsted, P., & Andersen, R. S. (2018). Cancer-before-cancer. *Medicine Anthropology Theory*, 5(5). https://doi.org/10.17157/mat.5.5.540

Pellegrino, E. D. (1994). Patient and physician autonomy: Conflicting rights and obligations in the physician-patient relationship. T*he Journal of Contemporary Health Law and Policy*, 10, 47–68.

Pols, J. (2015). Towards an empirical ethics in care: Relations with technologies in health care. *Medicine, Health Care and Philosophy,* 18(1), 81–90. https://doi.org/10.1007/s11019-014-9582-9

Region Hovedstaden. (n.d.). *Fordele og ulemper ved brystkræftscreening. Region Hovedstaden.* Retrieved May 7, 2025, from https://www.regionh.dk/Sundhed/Screening/Brystkr%C3%A6ftscreening/Sider/Fordele-og-ulemper.aspx

Region Hovedstaden. (2023). *Redegørelse om kapacitetsmæssige udfordringer på brystkræftområdet til Task Force for Patientforløb for Kræft- og Hjerteområdet.* Region Hovedstaden. https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerte/Bilag-2023_39-Region-Hovedstadens-redegoerelse_-brystkraeftomraadet.ashx

Region Midtjylland. (n.d.). *Screening for brystkræft.* Sundhed - Region Midtjylland. Retrieved December 13, 2024, from https://www.sundhed.rm.dk/sundhedstilbud/kraftscreening/screening-for-brystkraeft/

Region Midtjylland. (2023). *Redegørelse for udfordringer på brystkræftområdet.* Region Midtjylland. https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerte/Bilag/Bilag-2023_42-Region-Midtjyllands-redegoerelse_-brystkraeftomraadet.ashx

Robb, K. A., Simon, A. E., Miles, A., & Wardle, J. (2014). Public perceptions of cancer: A qualitative study of the balance of positive and negative beliefs. *BMJ Open*, 4(7), e005434. https://doi.org/10.1136/bmjopen-2014-005434

Schaarup, J. F. R., Aggarwal, R., Dalsgaard, E.-M., Norman, K., Dollerup, O. L., Ashrafian, H., Witte, D. R., Sandbæk, A., & Hulman, A. (2023). Perception of artificial intelligence-based solutions in healthcare among people with and without diabetes: A cross-sectional survey from the health in Central Denmark cohort. *Diabetes Epidemiology and Management*, 9, 100114. https://doi.org/10.1016/j.deman.2022.100114

Sheir, S., Manzini, A., Smith, H., & Ives, J. (2024). Adaptable robots, ethics, and trust: A qualitative and philosophical exploration of the individual experience of trustworthy AI. *AI & SOCIETY.* https://doi.org/10.1007/s00146-024-01938-8

Sundhedsstrukturkommissionen. (2024). S*undhedsstruktur kommissionens rapport—Beslutningsgrundlag for et mere lige, sammenhængende og bæredygtigt sundhedsvæsen.* Indenrigs- og Sundhedsministeriet. https://www.ism.dk/Media/638545635292256419/Hovedrapport-tilg%C3%A6ngelig-fil.pdf

The Boston Consulting Group. (2017). *Analyse af kapacitetsanvendelsen på kræftområdet.* https://www.ism.dk/Media/8/5/170407-kapacitetsanv-paa-kraeftomraadet.pdf

Van Cauwenberge, D., Van Biesen, W., Decruyenaere, J., Leune, T., & Sterckx, S. (2022). "Many roads lead to Rome and the Artificial Intelligence only shows me one road": An interview study on physician attitudes regarding the implementation of computerised clinical decision support systems. *BMC Medical Ethics*, 23(1), 50. https://doi.org/10.1186/s12910-022-00787-8

Wilhøft Kristensen, A., Lunde Jensen, A., Jensen, K., Oksbjerre Dalton, S., Friborg, J., & Grau, C. (2024). Exploring patient-reported barriers to participating in proton therapy clinical trials. *Technical Innovations & Patient Support in Radiation Oncology*, 29, 100230. https://doi.org/10.1016/j.tipsro.2023.100230

Witkowski, K., Okhai, R., & Neely, S. R. (2024). Public perceptions of artificial intelligence in healthcare: Ethical concerns and opportunities for patient-centered care. *BMC Medical Ethics*, 25(1), 74. https://doi.org/10.1186/s12910-024-01066-4

Zheng, J.-X., Li, X., Zhu, J., Guan, S.-Y., Zhang, S.-X., & Wang, W.-M. (2024). Interpretable machine learning for predicting chronic kidney disease progression risk. *DIGITAL HEALTH*, 10, 20552076231224225. https://doi.org/10.1177/20552076231224225