Huge Potential, Big Dilemmas

How to balance reporting on the AI revolution in healthcare

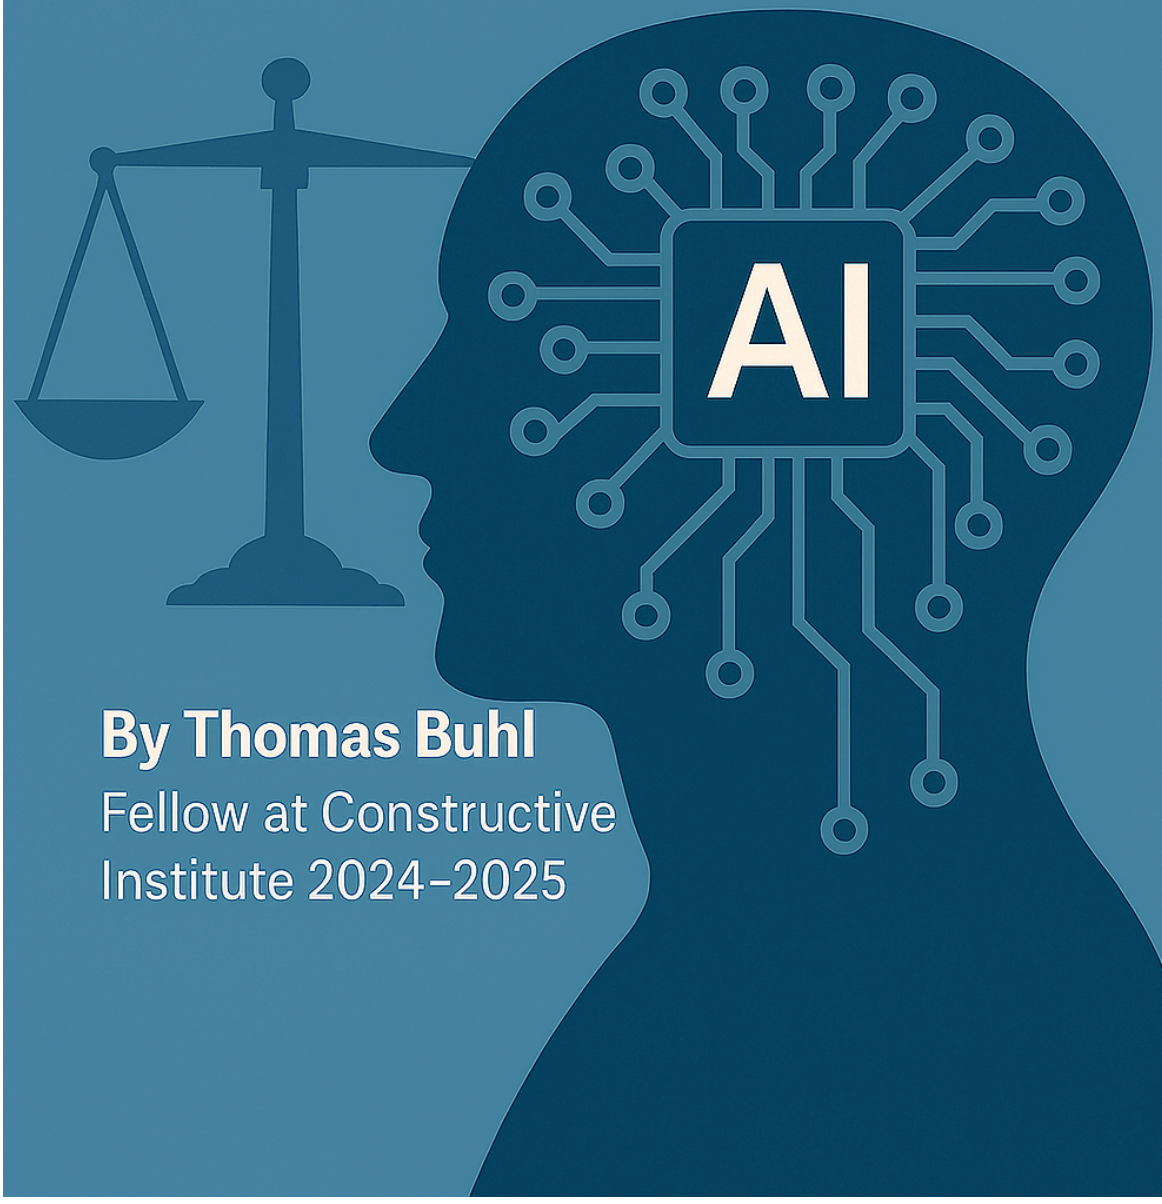

By Thomas Buhl
Fellow at Constructive Institute 2024–2025

Content

Resume

The media have a key role to play in conveying news and fostering a balanced, democratic conversation about AI in health – a technology that holds enormous potential but also poses significant dilemmas. This conversation is important because these very years, as a society, we are faced with the task of deciding how we want to use these new opportunities, and to what extent we should set boundaries through regulation.

This report shows that Danish mainstream media overwhelmingly cover AI in health with a clear focus on the promising potential, while dilemmas and challenges are mentioned only sporadically. At the same time, the chosen sources are predominantly AI experts, doctors, and industry professionals, while the patients – those whom AI is ultimately intended to serve – are virtually absent.

In other words, there is a lack of nuance and diversity – and of what constructive journalism regards as one of its core virtues: that the media 'see the world with both eyes'. In most cases, criticism of a one-sided worldview emerges when the media's negativity bias dominates, and stories of disasters, corruption, conflict, and scandal take centre stage. But it doesn’t seem much better to view reality uncritically only through unrealistically positive lenses.

Fortunately, there is a middle ground for the media – one that brings both the potential and the dilemmas of AI in health to light, helping ensure the balanced and informed public debate we deserve on such a vital and far-reaching topic. This report presents suggestions for how journalists and media outlets can strengthen their coverage and help bring this balance about.

Structure

Huge potential
The first part of the report explores how AI technologies are already being used in healthcare and touches upon emerging advancements and their potential impact on AI in healthcare in the coming years - from cancer diagnostics to brain-implanted microchips.

Big dilemmas
Secondly, the report identifies key ethical dilemmas associated with the implementation of AI in health. Topics covered include issues of accountability, algorithmic bias, lack of transparency in decision-making, and the broader societal implications of deploying AI technologies in healthcare settings.

Media coverage
The third section examines how the media cover AI in general, and AI in health in particular, in order to gain insight into how journalism shapes public understanding and perception. The centerpiece of this section is an analysis of 26 articles on AI in healthcare, published in Danish media over the course of 15 months. The findings reveal an overly positive framing of AI and a noticeable absence of patient and citizen perspective among the sources used.

Ideas and recommendations
The fourth and final section offers ideas and recommendations for how journalists can enhance their coverage of AI. Drawing on constructive journalism principles, it suggests ways to add nuance, engage the public and present balanced narratives that reflect both opportunities and challenges. Suggestions include critically examining proposed solutions, incorporating voices from diverse stakeholders and non-experts, and inviting readers and users into the conversation.

Illustration by ChatGPT. Prompt: Create an image that represents AI in healthcare

What I talk about when I talk about AI

So what exactly is artificial intelligence in healthcare? In this report it is based on a broad and general definition presented by Professor in Philosophy and AI researcher Jens Christian Bjerring from Aarhus University. He introduced it during a lecture on artificial intelligence in healthcare at the Folkeuniversitetet in Aarhus in February 2025:

Artificial intelligence in medicine centers on the use of machine learning algorithms that sift through vast amounts of medical data to uncover patterns or insights that can help improve health outcomes and patient experiences.

There are various branches and types of artificial intelligence, each with its own technical specifications and mechanisms. Some of these will be presented in the next chapter, which provides concrete examples of AI in healthcare. But the above - not exhaustive, yet relatively simple - definition serves as the basis for how the concept of AI in healthcare is understood and used in this report.

Huge potential

Ambitions for the use of AI in healthcare

The technology is no longer future speculation – it is already reality, and we see AI-assisted systems helping doctors with faster and more accurate diagnoses. Jens Christian Bjerring, Professor of Philosophy, Aarhus University, researcher of AI in health and decision-making processes.

In recent years, there has been an unprecedented surge in the development and application of artificial intelligence in healthcare. A clear testament to the growing market for AI health solutions is the explosive growth in AI-focused research studies within the field of medicine. According to data from Pubmed[1], Global research on AI in healthcare has grown exponentially. And The renowned science publication Nature[2] notes that AI is already revolutionizing healthcare.

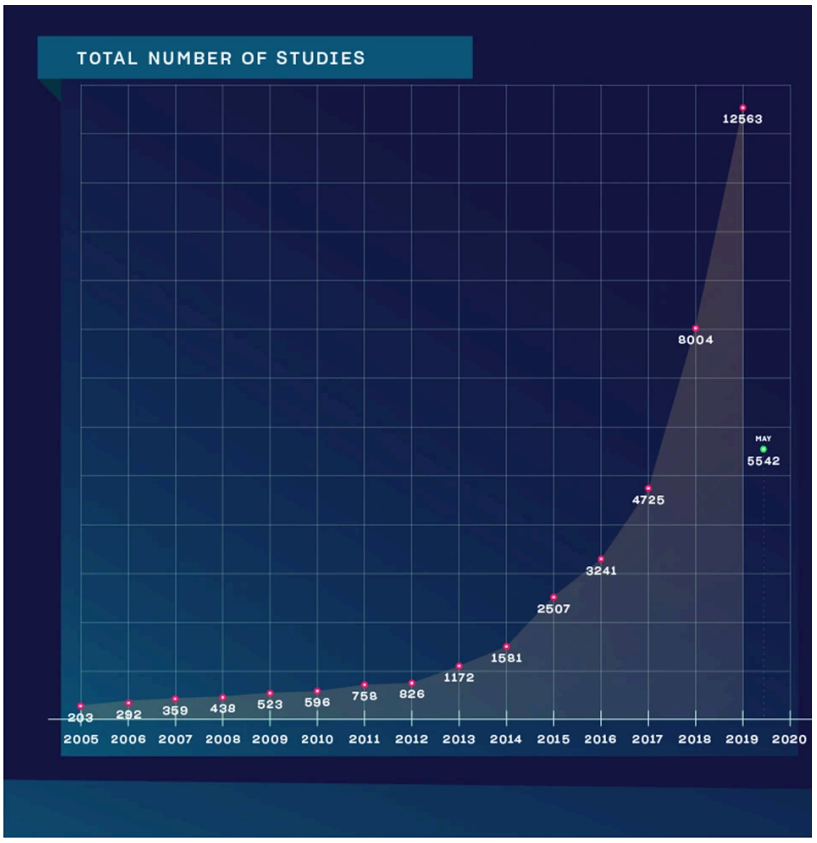

Number of global machine and deep learning studies year by year.*The last full year shown in the graph is 2019, when 12,563 studies were published. According to Pubmed, the number of studies has continued along the same exponential trajectory since then, with the most recent full year, 2024, setting a new record: 50,912 studies published. Source: pubmed.com*

At the national level, the Danish government has outlined an ambitious vision to become a global leader in the implementation of AI across public sectors, including healthcare. As it is

[1] 2020. With updated numbers 2020-2024. *Nature.*
[2] 2024. "AI in health care." *Nature.*

bluntly put in the government's AI-strategy[3] published December 2024: "Denmark aims at being world leader in the use of artificial intelligence in the public sector". This vision includes concrete measures to integrate AI technologies that improve efficiency and quality in the healthcare sector.

One of the goals of the strategy is to transfer tasks from medical staff to AI-systems. As Danish Minister of Health, Sophie Løhde, remarked in an interview in Berlingske[4] in 2023: "Artificial intelligence represents a significant opportunity to address the shortage of healthcare staff by automating repetitive tasks and optimizing resource allocation"

According to an investigative article in Politiken in December 2024[5], a dedicated government working group has already been tasked with estimating how many public sector jobs can be replaced through the use of AI. Politiken quotes written demands to the working group as stated in material that Politiken has had access to: "Real change on a large scale is demanded" as the working group is tasked with "identifying and implementing concrete solutions on a large scale that can free up at XX.XXX man-years by 2030". In other words: the objective for the working group is to find out how AI can replace the equivalent of between between 10.000 and 99.999 full time employees over the next 5-6 years

Against this backdrop of an exponentially growing body of AI-research, rapid technological progress and a political push to solve systemic problems in health care through artificial intelligence, the following section highlights current real-world examples of how AI is already supporting patient care and clinical workflows in Denmark - and where it's headed next.

Improved diagnostics

Image recognition and analysis are key strengths of artificial intelligence, helping doctors detect cancer in X-ray images. By training on thousands of annotated examples, AI systems learn to recognize patterns and identify subtle deviations from the norm — enabling early detection of disease.

In Denmark, this approach has been applied to breast cancer screening since 2021, when the Capital Region became the first region in the country[6] to allow AI to assist radiologists in making diagnoses. Radiologists retained full authority but were supported by AI as a diagnostic tool.

A 2024 study from the University of Copenhagen, based on more than 100,000 mammography images from the Capital Region's screening program, found that AI-assisted screening significantly improved outcomes: more cases of breast cancer were detected, and fewer false-positive results occurred. The use of AI also reduced the proportion of women who needed to be recalled for further examination by more than 20 percent. Additionally, the system proved so effective at identifying the 70 percent of images considered low risk that

[3] December 2024. "Strategisk indsats for kunstig intelligens - Et styrket fundament for ansvarlig udvikling og anvendelse af kunstig intelligens i Danmark." *Digitaliseringsministeriet.*
[4] Sidsel Lyhne Jensen. 15. august 2023. "Minister bebuder ai-offensiv: Skal afbøde mangel på varme hænder." *Berlingske.*
[5] Kristian Klarskov. December 2024. "Kæmpe AI-fremstød vil ændre tusinder, måske titusinder af offentligt ansattes hverdag fundamentalt." *Politiken.*
[6] Jeppe Tholstrup Bach. June 2024. Kunstig intelligens viser opløftende resultater i at opdage brystkræft. *tv2.dk*

hospitals were able to reduce the number of radiologists reviewing those cases from two to one.

Overall, the workload for radiologists was reduced by one-third, while diagnostic accuracy improved and more breast cancer cases were detected and treated in time. The method has since been adopted in several other Danish regions[7]

In a similar vein, a study publizied in The Lancet Oncology[8] showed that a radiologist working alongside an AI system was able to detect 20 percent more breast cancer cases in mammograms than two radiologists working together without AI support.

Detecting bone fracture

Bone fractures represent another area where artificial intelligence is being applied. In Denmark's northernmost region, Region Nordjylland, AI assists in identifying fractures by analyzing X-ray images. According to a press release from the region, the AI system is able to determine with 95 to 98 percent accuracy whether a fracture is present.

"We have found that AI performs on par with an experienced emergency room doctor," says Christian Pedersen, chief physician at Aalborg University Hospital, in an article at dr.dk.[9]

The AI assessments help physicians prioritize the most severe cases and determine appropriate treatment, leading to faster decision-making. As a result, emergency room wait times have been reduced by approximately one hour, and the number of patient complaints and compensation claims has dropped significantly. The method is now being rolled out to other parts of the country.

Predicting the course of disease

Artificial intelligence is also increasingly used to predict the course of disease for individual patients, with the goal of preventing complications, illness, and death through treatment options tailored to the individual's condition and medical history. This applies to both physical and mental health conditions, as the following two examples demonstrate.

At Zealand University Hospital in Køge, AI is transforming the treatment of patients with stomach or intestinal cancer by providing precise risk predictions. According to Region Zealand and Berlingske[10], the system is powered by a supercomputer trained on anonymized health data from 76,000 Danish patients who underwent surgery for colon or rectal cancer over the past decade. The AI integrates this data with other patient-specific information — such as age, sex, weight, blood test results, imaging data, comorbidities, prior

[7] Mathilde Bugge. June 2024."Hovedstaden får kunstig intelligens til at opdage brystkræft - resten af landet vil gøre den kunsten efter." *dr.dk*
[8] Kristina Lung et al. 2023. "Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial." *The Lancet*
[9] Allan Nisgaard og Simon Andersen Nielsen. February 2025 "Kan du finde fejlen på billedet? AI kan - og det kan få betydning for din næste tur på skadestuen" *dr.dk*
[10] Flemming Steen Pedersen. July 2023. "Tre uger efter stor kræftoperation var sportslegende tilbage på banen. Det kan han takke en computer for".*Berlingske*

diseases, risk factors, and overall health. Drawing from a dataset comprising 60,000 data points per patient, the system estimates the likelihood of complications or death within the first year after surgery.

Doctors then use these predictions to place patients into risk groups and implement tailored interventions. For example, vulnerable patients identified as high risk receive a preparatory regimen before surgery. This includes physical exercise programs, lung training to improve respiratory function, nutritional counseling, and iron supplements for those with anemia. These targeted efforts strengthen the most fragile patients before surgery, helping reduce the risk of serious post-operative complications.

Among the first 75 patients treated using this approach, none experienced severe post-surgical complications — compared to the usual 5 to 10 cases under conventional care. In addition, the time spent in hospital was halved for the most vulnerable patients, and readmissions were reduced to just two, compared to the usual 10 to 12. These results suggest that integrating AI into care not only improves survival but also enhances patients' quality of life and reduces pressure on hospital resources. The method is now being rolled out to other parts of the country, DR reported in maj 2025.[11]

Predicting risk of involuntary admission

Artificial intelligence is also being used in psychiatry to predict which patients are at the highest risk of becoming so ill that they require involuntary hospitalization. A study published online by Cambridge University Press[12] ,and conducted by researchers from Aarhus University and Region Midtjylland, showed that AI can help identify these patients months in advance.

"We believe this technology will help us intervene earlier, before patients reach a state where involuntary hospitalization becomes necessary," said Professor Søren Dinesen Østergaard in an article on Region Midtjyllands home page[13].

Søren Dinesen Østergaard, from the Department of Clinical Medicine at Aarhus University and Psychiatry in Region Midtjylland, contributed to the study, in which researchers trained AI on electronic health records from more than 50,000 voluntary psychiatric admissions in the region between 2013 and 2021. The system analyzed 1,800 variables per patient -including diagnoses, medication history, symptoms, prior use of coercion, lab results, and doctors' notes - to detect patterns indicating elevated risk of involuntary admission within six months of discharge.

For every 100 patients flagged as high risk by the algorithm, approximately 36 were in fact involuntarily hospitalized. Conversely, of every 100 patients assessed as low risk, 97 were

[11] August Olaf Jersild. 2025. "Kræftlægerne får en ny ven, og den er kunstig: 'Det er ikke sådan, at AI træffer en beslutning'." *dr.dk*
[12] Erik Perfalk et al. 2024. "Predicting involuntary admission following inpatient psychiatric treatment using machine learning trained on electronic health record data." *Cambridge University Press*
[13] Region Midtjylland. 2024. "Kan kunstig intelligens være med til at forbedre behandling af psykisk sygdom?" *Region Midtjylland*

not hospitalized. While the AI did not predict outcomes with perfect accuracy, the doctors involved in the project concluded that the system's level of precision is high enough to meaningfully assist in identifying patients in need of extra support — allowing earlier intervention and potentially preventing difficult, traumatic, or dangerous situations.

Personalizing treatment

In the two previous examples, AI was used to predict the course of disease and guide treatment plans tailored to the individual patient. More broadly, personalized treatment is a growing trend in healthcare. One that artificial intelligence has the potential to support and accelerate.

An example is found at Næstved Hospital, where AI is being used to personalize radiation therapy for cancer patients. In 2023, the hospital began using an advanced radiation therapy system known as *Ethos*, primarily for treating prostate and gynecological cancers. Traditional radiation therapy relies on static CT scans taken before treatment begins. In contrast, the Ethos system uses AI to adapt in real time to changes in the patient's anatomy.

Organs such as the bladder and intestines can shift in size and position between treatments depending on their contents. The Ethos system detects these changes and adjusts the radiation plan accordingly, ensuring that the tumor is targeted precisely while minimizing exposure to surrounding healthy tissue. This reduces the risk of side effects such as bowel or urinary issues.

Crucially, the technology allows doctors to increase the radiation dose without increasing side effects, thereby improving the likelihood of eliminating all cancer cells. This AI-driven approach not only improves treatment outcomes but also sets a new benchmark for individualized care in cancer therapy.

Saving time for general practitioners

Within general practice, artificial intelligence is already becoming a tool that helps practitioners (GPs) perform tasks more quickly and efficiently than before. At the same time, large language models such as ChatGPT may soon support doctors in making more accurate diagnoses.

According to an article published in the national Danish Daily Jyllands-Posten[14] in May 2025, a number of GPs in Denmark have already purchased a program with an AI assistant that listens in on consultations between doctor and patient.

*"*The program listens to the conversation between the doctor and the patient, instantly transcribes it in full, and then uses artificial intelligence to generate a complete medical note, which the doctor can review and insert into the patient's journal," said William

[14] Benjamin Jungdal. 2025. "Mere end hver 10. danske læge har købt en AI-assistent. Men det rejser en række alvorlige dilemmaer". *Jyllands-Posten*

Vossgård, CEO of the Norwegian company behind the solution, the AI firm Noteless, in an the article.

Although the solution raises ethical concerns - for instance, what happens if the AI makes errors in the medical notes - it has already been adopted by approximately 10 percent of Danish general practitioners, according to Jyllands-Posten. These doctors report that the technology helps reduce administrative burdens, potentially allowing them to spend more time with each patient.

Beyond handling paperwork at high speed and freeing up valuable time, AI may also assist general practitioners in making more precise clinical decisions. This applies even to large language models not specifically trained for healthcare use, according to Thomas Ploug, professor of data and AI ethics at Aalborg University, whom I interviewed at the university's Copenhagen campus.

"If we really want to make good use of ChatGPT, it might be in areas like identifying signs of diagnoses. Studies show that ChatGPT is incredibly effective at that", he says

A widely discussed study published in the Journal of the American Medical Association (JAMA)[15] and reported by New York Times[16] in November 2024 explored the impact of ChatGPT-4 on physicians' diagnostic accuracy. In a randomized clinical trial, 50 practicing physicians were divided into two groups: one with access to ChatGPT-4 in addition to conventional tools, and another using traditional resources only. All participants were asked to diagnose various cases based on written descriptions of patients and symptoms.

ChatGPT-4 was also tested independently — using the same patient descriptions, but without human input. The model alone provided the correct diagnosis in 90 percent of the cases. The doctors who were allowed to use ChatGPT achieved an average accuracy rate of 76 percent. Those who worked without AI reached an average of 74 percent.

The study's small sample size (only 50 doctors) and design raised questions. For example, the queries submitted to ChatGPT in its solo test were formulated by researchers, whereas doctors using ChatGPT had to formulate their own prompts. Despite these limitations, the study suggests that even general-purpose large language models like ChatGPT — not specifically developed for medical use — could and probably will be used by doctors in their work of identifying and naming physical ailments.

Saving lives by listening to phone calls

AI's full potential in healthcare is only beginning to unfold, with new opportunities and promises continuously emerging about how artificial intelligence can help save lives. One of the latest examples in Denmark comes from Region Zealand, which in 2025 became the first region in the country to launch a new AI tool aimed at reducing the long-term effects of strokes — including blood clots and brain hemorrhages.

---

[15] Ethan Goh, Robert Gallo, Jason Hom. 2024. "Large Language Model Influence on Diagnostic Reasoning." *JAMA Network*

[16] Gina Kolata. 2024. "A.I. Chatbots Defeated Doctors at Diagnosing Illness". *New York Times.*

An AI system now listens in on emergency 112 calls in the region, assisting dispatchers in identifying stroke symptoms more accurately and efficiently.

According to an article about the subject in Jyllands-Posten[17] strokes are the fourth leading cause of death in Denmark and the most common cause of adult disability. Patients who suffer a stroke require immediate treatment to avoid such consequences — every second counts. That's why it is crucial for emergency call operators to identify strokes quickly and accurately and dispatch an ambulance without delay.

Region Zealand implemented the AI system after a PhD project[18], in which doctor and researcher Jonathan Wenstrup trained the algorithm on 300,000 emergency calls from 2021. Among these calls, 757 were later confirmed as real strokes. Human operators — including nurses, paramedics, and physicians — had correctly identified 399 of them, or 52.7 percent. The AI system, in contrast, identified 477, raising the detection rate to 63 percent.

The AI also outperformed human operators in reducing false positives — cases where a stroke is suspected, but none is present. Human dispatchers misidentified 1,938 such cases, dispatching ambulances to patients who turned out not to have had a stroke. The AI system made fewer such errors, with 1,440 false positives. Reducing this number is critical, as ambulances tied up responding to false alarms may not be available for real emergencies.

In Region Zealand, emergency calls are still answered by human operators. But the AI listens in, and if it detects signs of a stroke based on how the conversation unfolds, it sends an alert to the dispatcher's screen. The operator then takes the final decision to send an ambulance or not.

The cyborgs are coming!

A new frontier in medical innovation is emerging at the intersection of artificial intelligence and neuroscience. One of the most spectacular developments in this field is the implantation of microchips in the human brain. For paralyzed patients, this technology could mean regaining the ability to move. For those with neurological disorders, it could offer new ways to communicate or access treatment. And in the future, the technology may be used with the aim of enhancing our cognitive and intellectual capacities. Beyond its medical implications, this development also raises profound ethical and philosophical questions. What does it mean to be human when technology becomes an integral part of our cognitive functions? How can we ensure that such innovations benefit individuals and society without compromising autonomy or privacy or creating extreme inequity?

One of the most widely publicized brain-implant experiments is being conducted by Neuralink, a company owned by Elon Musk. In August 2024, Neuralinkfor the second time implanted its device, known as The Link, into the brain of a patient[19], Alex, who is paralyzed from the neck down due to a spinal cord injury. The implant, roughly the size of a Danish

[17] Thomas Møller Kristensen. 2025. "Ny metode skal hjælpe med at afsløre, om dit alarmopkald gælder en af Danmarks allerdyreste sygdomme." *Jyllands-Posten*

[18] Jonathan Wenstrup. 2024. "Telephone dispatcher recognition of stroke and machine learning as a possible supportive tool." *Region Hovedstaden*

[19] Rachael Levy. 2024. " Neuralink implanted second trial patient with brain chip, Musk says." *Reuters*

five-kroner coin, consists of ultra-thin electrodes that extend into the brain's neurons and capture the electrical signals generated when a person thinks. These signals are wirelessly transmitted to a computer, where an AI system filters and interprets them, translating thoughts into actions.

As the AI system trains, it learns to recognize patterns in brain activity and associate them with the actions the patient intends to perform, such as moving a computer cursor or operating a prosthetic device.

Using this technology, Alex has been able to play the video game *Counter-Strike 2* using only the power of thought. Similarly, according to Neuralink, he has also used 3D design software by thought alone.

More breakthroughs have followed in recent years. In 2024, the British scientific journal Nature published the study Walking Naturally After Spinal Cord Injury Using a Brain–Spine Interface[20] The study documented how a paralyzed patient regained the ability to walk using brain-implanted electrodes. Similar to the Neuralink case, the electrodes recorded brain signals produced when the patient thought about movement. Artificial intelligence then processed the signals and wirelessly transmitted them to a stimulator implanted in the lower spine. The stimulator activated the necessary muscles, enabling the patient to walk, climb stairs, and navigate uneven ground.

This technology still requires further development before it can be widely deployed to help people with paralysis. But even that may be just the beginning. Neuralink is also developing *Project Blindsight*, an initiative that aims to restore partial vision to blind individuals using brain-implanted microchips.

In the long run, Elon Musk envisions a future in which brain chips allow ordinary people to interact with computers through thought alone, enabling seamless communication with AI. He sees this as a way to enhance human capabilities[21], and protect humanity from the risk of artificial intelligence taking over.

This vision of humans moving further down what could be called a "cyborg path" in order to combat AI domination inevitably raises fundamental ethical and philosophical dilemmas - questions about human identity, embodiment, and our place in the world.

But wait - there's more!

This section's overview of the use of artificial intelligence in healthcare only scratches the surface of the many ways AI is currently being applied, and how it will be implemented in the near future.

I could also have mentioned how AI is accelerating drug and vaccine development, with algorithms capable of predicting efficacy and safety far more quickly than traditional methods. AI is also being used to screen molecular compounds, speeding up the identification of potential treatments for cancer and other life-threatening diseases.

---

[20] Henri Lorach et al. 2023. "Walking naturally after spinal cord injury using a brain–spine interface." *Nature*
[21] Robert Hart. 2023."Elon Musk Says Neuralink Could Slash Risk From AI As Firm Prepares For First Human Trials." *Forbes*

AI algorithms further have the potential for being used for remote patient monitoring by accessing health data, heart rate, and other key indicators captured by patients' smartphones, enabling early intervention and reducing the risk of hospital readmissions. Naturally, this raises important questions. For instance: What level of individual surveillance are we willing to accept?

A similar perspective - and a similar dilemma - was raised by Associate Professor of Psychology Malene Flensborg Damholdt from Aarhus University during a lecture at Folkeuniversitetet in Aarhus in February 2025. She pointed out that AI may very well one day assist in diagnosing mental illnesses by monitoring and analyzing a patient's social media activity, smartwatch data, health apps, and other digital biomarkers.

Some researchers even predict that AI will be able to analyze medical records and, based on the frequency of certain keywords or descriptions, forecast the risk of hospitalization for elderly patients living at home. Before care staff detect any issues.

Having now touched upon a few of the practical and potential uses of AI in healthcare, amid its exponential growth, this report now turns to the practical, ethical, and philosophical dilemmas that follow. These will be explored in the next section.

Big dilemmas

The bias dilemma

A fundamental ethical concern in the implementation of artificial intelligence in healthcare is the risk of reinforcing existing biases in medical data, with potentially serious consequences for patients.

Jens Christian Bjerring, Professor of philosophy at Aarhus University and researcher of AI in healthcare, explains the problem through two examples during an interview. Historically, more research has been conducted on men than women, and on white populations more than people of color. As a result, health datasets contain disproportionate amounts of information about men and white individuals.

“When we train algorithms on these kinds of datasets, the risk is that the algorithms become better at making predictions for men than for women — and better for white people than for black people”, says Jens Christian Bjerring.

This means that bias embedded in the training data can be reproduced and amplified in the AI models deployed in healthcare systems — influencing predictions of disease trajectories and treatment recommendations.

Bias in skin cancer diagnosis

This issue has already emerged in AI models designed to detect skin cancer by analyzing images of skin irregularities. A 2021 study[22] showed that many such models were trained primarily on images of light-skinned individuals. Another scientific investigation[23] found that these models’ diagnostic accuracy dropped by 27–36 percent when tested on a dataset that included patients with diverse skin tones. This significantly increases the risk that skin cancer in darker-skinned patients will be diagnosed at a later stage — potentially reducing survival rates

Attempts at predicting depression

A 2023 case study from the University of Copenhagen[24] demonstrated how systematic bias had become embedded in an AI model developed to predict the risk of severe depression. The model was trained on national registry data from 240,000 individuals, incorporating variables such as age, gender, income, education level, marital status, residential area, number of address changes, immigration status, and many others.

---

[22] Matthew Groh et al. 2021. “Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset.” *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*

[23] Elizabeth Bailey et al. 2022. “Disparities in dermatology AI performance on a diverse, curated clinical image set” *Science*

[24] Eike Petersen et al. 2022. “On (assessing) the fairness of risk score models.” *Association for Computing Machinery*

The study found that the model was better at predicting depression in women than in men. It gave less accurate risk assessments for low-income individuals than for high-income individuals, performed better for urban residents than for those in rural areas, and showed significantly lower precision for non-western immigrants compared to the ethnic majority. It also systematically underestimated the risk of depression among younger people.

Researchers warned[25] that such biased predictions could result in unequal access to care if used to determine who receives preventive treatment or how healthcare resources are allocated. They called for more diverse training data from underrepresented groups and improved model calibration across demographic categories before deploying such systems in practice.

Bias is subtle and often hidden

In July 2024, three researchers published an article in Ugeskrift for Læger[26] warning that: "AI tools may unintentionally reproduce and amplify inequalities in the healthcare system."

They concluded that "Even with diligence and the best intentions, it can be difficult to determine in advance how skewed models can be avoided."

While AI models can be recalibrated, adjusted, or retrained when bias is discovered, the authors emphasized that such distortions are hard to eliminate entirely.

Professor Thomas Ploug, a leading expert on data and AI ethics at Aalborg University, author of *Medical Ethics*, and former member of the Danish Council on Ethics (2010–2016), confirms that new AI models are tested for bias — sometimes through extensive procedures. Yet detecting where and how bias originates remains a challenge.

"It's often very subtle. It can be difficult to know what exactly in the training data is causing the model to make certain recommendations", Ploug says.

According to the professor, the challenge becomes especially serious when AI systems reach their conclusions through complex and opaque processes that even developers cannot fully explain.

This lack of transparency leads to one of the most pressing dilemmas in AI healthcare: the black box problem. When the logic behind an algorithm's recommendation is inaccessible, even to its creators, it becomes difficult to understand, challenge, or regulate the decisions it produces. And in the context of healthcare, where decisions can have life-altering consequences, that uncertainty poses a serious ethical concern.

---

[25] Kristian Bjørn Hansen. 2023. "Blind brug af AI i sundhedsvæsnet kan skabe usynlig forskelsbehandling." *Det natur- og biovidenskabelige fakultet, Københavns Universitet*

[26] Anne Gerdes et al. 2024. "AI-etik i sundhedsvæsenet." *Ugeskrift for læger*

The black box dilemma

AI machine learning models — particularly deep neural networks — are often extremely difficult to understand, even for experts and physicians. This creates challenges for healthcare professionals, who must be able to explain diagnostic and treatment decisions to patients.

Many AI systems used in healthcare are based on such deep neural networks that are fed complex input data — often containing thousands of data points about a given patient. These may include information on prior illnesses, weight, gender, genetic traits, social background, and many more. The AI model processes these data through millions of mathematical functions to identify patterns, which are then compared to the training cases it has previously seen. Based on this, the system generates an output. For example, a probability that the patient suffers from a particular condition.

Because the system performs such a vast number of internal operations to reach its conclusion, it is often difficult, if not impossible, for medical staff to determine how and why the system reached a specific result. The calculations take place, figuratively speaking, inside a sealed black box to which we have no access. This is the black box problem.

“We can observe that an AI system is good at producing input-output pairings with high precision, but if we want to understand how it arrived at a specific output, we often have no idea how it got there,” says Jens Christian Bjerring, Professor of philosophy and researcher in AI and healthcare at Aarhus University.

This lack of transparency can lead to skepticism from both physicians and patients toward AI-generated decisions and recommendations: How can a doctor trust a diagnosis without knowing the rationale behind it? And how can a patient make an informed decision to accept an AI-recommended treatment when no meaningful explanation is provided?

“The lack of explainability in complex AI models creates serious challenges for trust and accountability in clinical decision-making,” says Thomas Ploug, professor of data and AI ethics at Aalborg University.

Predicting psychosis through language

As an early example of the black box phenomenon in healthcare, Ploug refers to a 2015 study published in Nature[27]. In the study, researchers explored how automated language analysis might be used to predict the development of psychosis among young people who had already shown early warning signs.

The researchers conducted detailed interviews with 34 at risk individuals and then allowed AI language models to analyze the transcripts. The AI system correctly identified - with 100 percent accuracy - the five individuals who later developed psychosis. Among the linguistic

[27] Gillinder Bevi et al. 2015. “Automated analysis of free speech predicts psychosis onset in high-risk youths.” *Nature*

markers the AI highlighted was a reduced use of the words “which” and “that.” Other patterns included shorter sentences and weaker semantic coherence between sentences.

But how can using the word “which” less frequently be a predictor of psychosis? The AI model found a statistical correlation. But not an explanation. The reasoning behind the association was lost inside the black box. Understanding why this linguistic signal might be relevant would require extensive follow-up research in linguistics and psychology.

“This is one of the fundamental issues with artificial intelligence,” says Ploug. “It identifies correlations, patterns. But that doesn’t mean there’s causality. We can't produce a causal explanation model that tells us why reduced use of the word ‘which’ would indicate a higher risk of psychosis.”

A serious barrier to implementing AI in health

According to Jens Christian Bjerring, the black box dilemma is one of the most serious barriers to implementing AI in diagnostic and treatment settings. He offers a hypothetical example to illustrate:

“Imagine an AI model that is statistically more accurate at predicting testicular cancer than human doctors. A physician inputs all your symptoms and health data into the model. The AI calculates that you have a very high risk of developing testicular cancer and recommends preemptive removal of your testicles. Naturally, before agreeing, you want to know why the system reached that conclusion. But because the AI model is a black box, the physician may not know what data or patterns the model used to arrive at its recommendation.”

Even if the system is statistically more accurate than human doctors, no AI model is 100 percent correct. They also make mistakes. So, can we trust a machine with such a serious decision if we don’t understand how it arrived at it?

Responsibility, ethics, and informed consent

A 2024 article by three Danish doctors and researchers[28] in the medical magazine “Ugeskrift for Læger” discusses a similar issue based on a study published in the Journal of Medical Internet Research[29]. In that study, an opaque but diagnostically validated AI system ranks multiple treatment options for a breast cancer patient and recommends the surgical removal of both breasts. The Danish authors argue that knowing the system is validated is not enough. Without additional information about how the model reached its recommendation, neither the physician nor the patient can make a responsible, well-informed treatment decision.

[28] Anne Gerdes et al. 2024. “AI-etik i sundhedsvæsenet.” *Ugeskrift for læger*

[29] Jens Christian Bjerring og Jacob Busch. 2020. “Artificial Intelligence and Patient-Centered Decision-Making.” *Philos.Technol*

This scenario, says Bjerring, highlights another variant of the black box dilemma: If AI systems cannot provide a clear and understandable rationale for their decisions, they may come into conflict with legal and ethical framework, including Danish healthcare law. Section 16 of the Danish Health Act states that "the patient has the right to receive information about his or her health condition", and that such information "must be provided on an ongoing basis and provide an understandable presentation of the disease" in order to enable informed consent to treatment.

"But when an AI system cannot offer a coherent explanation," says Bjerring, "that informed consent becomes difficult, if not impossible. It is an explainability dilemma that so far has been very hard to solve.

The explainability dilemma

So far, attempts at having one AI analyse and explain what happens inside another AI when it reaches its "black box" decisions, have not been particularly fruitful. That method, called Explainable AI, or XAI for short, could otherwise be a way to provide patients with understandable explanations for diagnoses or treatment recommendations, thereby supporting the principle of informed consent as required by healthcare law.

"There are many types of XAI, but so far, their effectiveness in real-world settings remains a matter of ongoing debate" ,says Jens Christian Bjerring, philosopher and AI-in-healthcare researcher at Aarhus University.

So the dilemma remains. Because explainability is a high priority for the patient. A study conducted in 2021 by, among others, Thomas Ploug, professor of data and AI ethics at Aalborg University, and published in the Journal of Medical Internet Research[30], asked 1,072 Danish respondents about their attitudes toward AI-supported diagnoses and treatment recommendations.

The participants' top concern was that it should always be the doctor, not the AI, who bears final responsibility for diagnosis and treatment planning. Their second-highest priority was that any AI used in clinical settings must be explainable.

Anne Gerdes, professor and AI ethics researcher at the University of Southern Denmark, distinguishes between high-risk and low-risk AI decisions in a 2024 article in Springer Nature[31]. For low-risk situations such as prioritizing patients for knee scans, limited explainability may be acceptable. But in high-risk contexts such as cancer treatment, more comprehensive explanations are necessary.

This is also reflected in the EU's AI Act, which stipulates: "High-risk AI systems shall be designed and developed in such a way to ensure that their operation is sufficiently transparent to enable users to interpret the system's output and use it appropriately." That challenge remains unsolved in practice.

---

[30] Thomas Ploug et al. 2021."Population Preferences for Performance and Explainability of Artificial Intelligence in Health Care: Choice-Based Conjoint Survey" *Journal of Medical Internet Research*

[31] Anne Gerdes. 2024, "The role of explainability in AI-supported medical decision-making"*Springer Nature*

All in all, a lack of transparency, limited explainability, and medical decisions based on non-transparent AI analyses give rise to a new ethical concern: the responsibility dilemma.

The responsibility dilemma

The use of artificial intelligence in the healthcare sector raises a pressing dilemma about responsibility. Especially when AI systems contribute to serious errors in diagnosis or treatment. Traditionally, responsibility has rested with healthcare professionals, but when AI technology plays a central role in decision-making, it becomes less clear who should be held accountable: the software developers, the healthcare professionals, the AI manufacturer, or the healthcare institution itself.

Jens Christian Bjerring, philosopher and AI researcher, illustrates the problem with a thought experiment:

“Imagine a doctor who knows that the AI he uses to help diagnose patients is, on average, 10 percent more accurate than he is. Now imagine that this AI system recommends a woman have one of her breasts removed due to a cancer risk. Even though the AI has reached this recommendation through an opaque process, and the doctor doesn’t fully understand how the AI arrived at its conclusion, he still has strong reason to follow the recommendation, because he knows the AI system is generally more precise. But no AI system is infallible. What happens if the diagnosis is wrong? Who is responsible?” asks Jens Christian Bjerring.

“If the doctor cannot explain how the algorithm arrived at the prediction on which the medical decision is based, then we are dealing with what the literature calls a *responsibility gap”*, he explains.

“In practice, the doctor would carry the responsibility, but it’s not necessarily comfortable for doctors to use this kind of technology if they are ultimately accountable for something they don’t understand”, says Bjerring.

On the other hand, it might also be problematic - or even legally risky - not to follow an AI recommendation when the algorithm demonstrably has a higher diagnostic success rate than the doctor.

Healthcare professionals in Region Zealand have sought to guard themselves against such liability in the previously mentioned project, where an AI listens in on emergency calls to 112 and sends an alert to the operator’s screen if it detects signs of a stroke. According to Jyllands-Posten[32], the emergency call handlers requested legal assurances that they would not be held accountable if they chose to disregard the system’s alert in a case where a patient later turned out to have suffered a stroke. Those guarantees have been provided:

“They have no reason to fear repercussions, as long as they can offer a medically sound

[32] Thomas Møller Kristensen. 2025. “Ny metode skal hjælpe med at afsløre, om dit alarmopkald gælder en af Danmarks allerdyreste sygdomme.” *Jyllands-Posten*

justification for their decisions," says Nikolaj Blomberg, senior consultant in the region's IT department, to Jyllands-Posten.

Danish AI researchers and physicians Anne Gerdes, Iben Fasterholdt, and Benjamin S. B. Rasmussen argued in a 2024 article in Ugeskrift for Læger[33] that AI in healthcare must not be too opaque: "When AI tools contribute input to decision-making, it is a matter of essential legitimacy that the AI's role is transparent and that responsibility can be clearly placed."

They also acknowledge, however, that some believe performance should take precedence over explainability and over the ability to assign clear responsibility when mistakes occur: "Pushed to its logical conclusion, the argument goes that saving lives is more important than being able to explain how a black-box model arrived at its recommendation," they write.

From a legal standpoint, responsibility for AI-supported decisions varies across countries. But there seems to be a clear need for transparent frameworks and regulation. Without precise guidelines, trust in AI-powered healthcare systems may erode, potentially slowing down adoption of the technology, even where it could bring clear benefits.

The question remains whether there is enough time to establish the necessary regulations and protocols. There is mounting political pressure from government and parliament to rapidly accelerate AI implementation in the healthcare sector. And this raises a critical concern: Are we at risk of pushing flawed AI systems into practice too quickly, before they've been adequately validated and tested?

The speed/safety dilemma

"We face a complex dilemma where there are opposing considerations between pressure for rapid implementation of AI in healthcare and the need for thorough development and testing of AI tools before deployment."

This is how Anne Gerdes, Professor and visiting researcher at the Center for Clinical AI at Odense University Hospital (OUH), senior researcher Iben Fasterholdt from the Center for Innovative Medical Technology (OUH), and physician PhD Benjamin S. B. Rasmussen from the Department of Radiology (OUH) frames the issue in their 2024 article of the Danish medical periodical, Ugeskrift for Læger[34].

AI development is moving so fast that it will have consequences, and not only joyful ones, according to Thomas Ploug, Professor of Data and AI Ethics at Aalborg University and author of the book Medical Ethics.

"Technology is evolving and being implemented so quickly that errors are inevitable. And we will see AI introduced in areas where it never should have been deployed," he says..

Ploug draws parallels to other public sector implementations, noting:

"If we look at the short history of implementing AI in the Danish public sector, there have

[33] Anne Gerdes et al. 2024. "AI-etik i sundhedsvæsenet." *Ugeskrift for læger*
[34] Anne Gerdes et al. 2024. "AI-etik i sundhedsvæsenet." *Ugeskrift for læger*

already been several scandals. Just look at the property valuation system. There's no reason to believe similar issues won't happen again, including in healthcare."

As previously mentioned, the Danish government has announced its ambition to become a global leader in the implementation of AI across all public sectors. And according to the Danish Minister of Health, Sophie Løhde, AI in health needs to be implemented as soon as possible, in order to solve pressing problems in Danish healthcare.

"Artificial intelligence represents a significant opportunity to address the shortage of healthcare staff by automating repetitive tasks and optimizing resource allocation", she stated in an interview with Berlingske in 2023[35].

And according to Politiken in December 2024[36], a dedicated government task force has already been assigned to determine how many public sector jobs could be replaced by AI.

In short, the pressure to implement AI in healthcare is intensifying. But the rapid pace of this rollout presents challenges. One of the greatest risks is that AI tools may not be sufficiently validated in clinical, Danish settings before use, potentially leading to misdiagnosis or inappropriate treatment recommendations, with real consequences for patient safety.

Moreover, fast-tracking AI without thorough ethical review or regulatory frameworks increases the risk of unclear lines of accountability or even legal liability, if an AI system causes harm.

As Gerdes, Fasterholdt, and Rasmussen emphasize in the Ugeskrift for læger-article:

There is a conflict between the pressure to implement AI in healthcare and the need for comprehensive, time-consuming empirical validation of these technologies. Therefore, it must be emphasized that validating AI models requires careful evaluations and, during implementation of AI, a strict control regime.

Other challenges

The overtreatment dilemma

According to Danish Minister of Health, Sophie Løhde, artificial intelligence is meant to help address some of the healthcare system's capacity issues. But this promise comes with an inherent dilemma: AI enables the identification of an increasing number of individual health risk parameters, which may lead to overdiagnosis and pressure toward overtreatment.

This concern was raised in an op-ed by Christoffer Bjerre Haase, physician and researcher at the University of Copenhagen, published in Information in February 2025[37] He asks: What

[35] Sidsel Lyhne Jensen. 2023. "Minister bebuder ai-offensiv: Skal afbøde mangel på varme hænder." *Berlingske*
[36] Kristian Klarskov. 2024. "Kæmpe AI-fremstød vil ændre tusinder, måske titusinder af offentligt ansattes hverdag fundamentalt." *Politiken*
[37] Christoffer Bjerre Haase. 2025."Techgiganternes magt skyller ind over sundhedsområdet som en tsunami, og vi er uforberedte." *Information*

do we do when AI systems, via smartwatches, full-body scans, or self-tests, identify dozens of potential health risks for a single individual?

“If AI finds you have a five percent or even a one percent risk, does it benefit society to treat you? If we do, there will be a lot of people who with 99 percent certainty are treated unnecessarily,” Haase writes, and continues:

“The healthcare system is already under strain. Acting on every potential health risk could stretch resources to the point where patients with acute, psychological, or social needs are deprioritized in favor of those with risk-based diagnoses.” And that, Haase writes, “is a delicate dilemma.”

The regulation dilemma

The rapid AI development risks outpacing the laws and safeguards designed to protect citizens. In March 2024, the European Parliament adopted the AI Act—described as the world’s first comprehensive legislation on artificial intelligence. Its aim is to establish clear frameworks for developers, distributors, and users of AI, striking a balance between corporate innovation and citizens’ rights.

The Act is being phased in between 2025 and 2027, but it’s already under pressure from start-ups and large tech companies who argue that the legislation imposes excessive administrative burdens. Others warn that, while the EU focuses on regulation, it risks falling behind the U.S. and China in technological development.

Meanwhile, the European Commission is still working to finalize the detailed regulations and guidelines that were postponed when the Act was passed.

Domestically, philosopher and AI researcher Jens Christian Bjerring doubts that legislation can keep up:

“I’m not optimistic that we can respond to technological developments quickly enough,” he says. “How are we supposed to legislate for technologies we don’t even know will exist two years from now? It’s incredibly difficult.”

The cyborg dilemma

The implantation of microchips in the human brain, touched on in the previous chapter, raises not only practical and legal challenges, but deeper philosophical ones.

If Elon Musk’s vision becomes reality, brain chips will not only help blind people see and paralyzed individuals walk. They may also be used to enhance the cognitive abilities of healthy people by fusing human brains with AI.

Jens Christian Bjerring points to a range of challenges:
“If more of us begin having chips implanted in our heads, directly connected to AI systems, what happens to data security? What if someone hacks the biochip in your brain? What

about the right to privacy—who should be allowed access to your brain data?" he asks. "These are the kinds of ethical problems we'll have to confront."

At stake are not only rights and responsibilities but fundamental questions about what it means to be human in a future where mind and machine converge.

I have touched upon some of the dilemmas, AI in health pose. Several additional questions could deserve attention: What does our increasing reliance on health technologies mean for the relationship between healthcare professionals and patients: How do we safeguard against dehumanized care, where human contact and compassion are sidelined? How do we protect the right to privacy as more and more of our health data becomes accessible and used by AI systems?

In short, there is no shortage of dilemmas. The purpose of the previous two chapters has been to show that AI in healthcare holds vast potential, *and* presents serious and complex challenges.

If citizens are to form informed opinions about this and contribute meaningfully to a democratic debate that can guide political decisions, *both* opportunities *and* dilemmas must be placed on the public agenda. Here, media coverage plays a critical role.

So how do the media cover AI in healthcare? These are the themes explored in the remainder of this report.

How Danish media covers AI in healthcare

Are mainstream media helping to fuel the hype surrounding artificial intelligence and its vast potential in healthcare? Do they offer a nuanced perspective that includes both opportunities and challenges? Or is the coverage predominantly focused on critical concerns?

To shed light on these questions, and to get a clearer picture of whose voices are included when the media report on AI in healthcare, I analyzed selected articles published in Danish mainstream outlets between January 1, 2024, and April 1, 2025.

The aim of this analysis is to contribute to a broader conversation about whether the media are fulfilling their role as facilitators of informed and balanced public debate on a vital societal development, or whether there is room for improvement.

Danish media embrace AI in healthcare

The findings are striking: Of the 26 articles examined that focus on AI in healthcare, 22 present the topic in an overwhelmingly positive light. Only two articles adopt a more balanced approach, weighing both the benefits and potential drawbacks. And just two articles have a predominantly critical framing, drawing attention to the ethical and practical challenges that come with implementing AI in the health sector. *(A full overview of all 26 articles is available in the appendix)*

Looking at the sources quoted, the majority are domain experts, primarily researchers and physicians with specialist knowledge of AI, along with business leaders and political figures. Out of a total of 53 sources cited, only one belongs to the group of people whose lives and health are directly affected by how AI is applied: the patients. There are only three other instances where an organizations representing patients' interest is cited as a source, namely the Danish Cancer Society (Kræftens Bekæmpelse). Furthermore, not a single source with professional expertise in data ethics or digital privacy is represented. *(All sources are listed in the appendix.)*

How I did it

The analysis is based on a search through the Infomedia database, covering all Danish-language national, regional, and local newspapers, including their websites, published between January 1, 2024, and March 31, 2025. The search terms used were "AI" and "sundhed*" (health*), meaning that only articles containing both the term "AI" and words beginning with "sundhed" (such as sundhed (health), sundhedsvæsen (health care), or sundhedssystemet (health care system)) were included.

The initial search produced 419 articles.

I then excluded all articles that, despite matching the search terms, clearly focused on other topics and mentioned AI and health only in passing.

Duplicate stories, typically identical Ritzau newswire pieces reprinted across several outlets, were only counted once.

Finally, I excluded non-editorial content, such as op-eds, letters to the editor, and opinion pieces in order to focus on the content which was a direct result of editorial decisions about which stories and angles the media itself had prepared.

This filtering left 26 original news articles that explicitly addressed the use of AI in healthcare. Each article was read thoroughly and categorized according to its overall framing: positive, neutral, or critical. This was an easier task than one might imagine, since the language used, the quoted sources, and the overall tone of the vast majority of the articles left no doubt as to which category they belonged to. All quoted sources were categorized into six groups: experts, business representatives, politicians, civil servants, NGOs/interest groups, and patients.

Articles in total	26
Positively angled articles	22
Neutrally angled articles	2
Critically angled articles	2
Total number of sources	53
Experts (researchers, doctors, AI-experts)	24
Business representatives	8
Politicians	8
Civil servants	6
NGO/Interest organizations	5
Patients	1

Study of Danish-language newspapers' coverage of AI in health in the period 1/1 2024 to 1/4 2025.

A Strong positivity bias

Out of 26 articles present an unequivocally positive framing of AI in healthcare. Most of these articles highlight the transformative potential of AI in improving diagnostics, treatment, and efficiency in healthcare. They carry headlines such as:

- *Artificial intelligence to help men get faster cancer results* (Nordjyske Stiftstidende, 4/2 202522)

- *Artificial intelligence can predict who is at high risk of involuntary hospitalization* (Jyllands-Posten, 28/11 2024)

- *Computer detects cancer, cardiac arrest and broken bones. AI surges ahead in hospitals* (Berlingske, 22/2 2024)

- *Huge potential* (Ekstra Bladet, 2/11 2024).

Others of the positive articles highlight systemic or regulatory barriers that prevent AI from being more widely implemented to the benefit of patients. While these pieces adopt a critical tone, the criticism is not directed at AI. On the contrary. They carry headlines such as

- *Lack of legislation blocks the use of artificial intelligence in elderly care* (Berlingske, 3/3 2024)
- *Danish AI companies hit a wall with public sector lawyers* (Jyllands-Posten, 28/3 2025).

Two articles take a more balanced, neutral stance:

- *Researchers want to prevent suicide with AI* (Flensborg Avis, 24/8 2024)

- *SF and PFA in historic alliance: AI to make Danes healthie*" (Berlingske, 8/3 2025)

Both of these articles explore both the promise of AI, and also its potential risks and dilemmas.

Only two of the 26 articles have a predominantly critical angle. One of them is

- *Hidden behind the screen* (Weekendavisen, 6/9 2024)

And it includes the only patient source in the entire set of 53 cited sources. It addresses the use of AI in psychiatry, where researchers are developing models that analyze patient records to assess the risk of readmission. The featured patient, who has previously been hospitalized with severe depression, expresses deep skepticism, fearing that such systems could dehumanize care and shift the doctor's attention away from the person and toward data-driven predictions.

The other critical piece is

- *Meant to save the municipality millions – large AI project must be scrapped* (Frederiksborg Amts Avis, 22/7 2024)

It reports on a failed AI implementation in Gribskov Municipality. The supplier could not get the AI system to function as expected, and the project was abandoned. A stark contrast to the same media's earlier coverage highlighting the very same AI systems promised savings.

Taken together, these findings suggest a media landscape largely enthusiastic about AI in healthcare. Ethical dilemmas, implementation challenges and long-term risks are only sporadically represented. And those who are ultimately affected, the patients, have almost no voice. While many individual articles are informative and well-crafted, the overall coverage shows limited nuance and a consistent heavy leaning toward the positive.

Limitations of the study

The strength of the conclusions in this study can be discussed, as the study has a number of limitations. One of them is that it includes a relatively modest number of articles.

Moreover, the Infomedia search only included the words “AI” and “health*”. One could have supplemented with other search strings, and that could likely have produced a larger number of relevant articles and thereby created a more solid data foundation. For example, one could have searched for “AI and hospitals”, “artificial intelligence and health”, or “artificial intelligence and treatment”, just to name a few of the alternative possibilities.

It can also be discussed whether the limitation to only focus on the editorial content gives an accurate picture of the media’s output. The excluded letters to the editor, columns, op-eds, and so on can also be seen as media output, and it could potentially have nuanced the picture to include them in the analysis.

Finally, the limitation to national, regional, and local newspapers does not represent the full media landscape. If I had also chosen to search TV and radio segments, web sources, magazines, etc., it might have given a more complete picture.

The study should therefore be assessed in light of the choices I have made in order to make the study feasible within the time frame I had available. Despite these limitations, the consistency of the findings suggests clear tendencies worth further exploration. The study can be seen as a first attempt to map tendencies in Danish media coverage of AI in healthcare.

International perspective

To my knowledge, no other studies have specifically examined how Danish media cover AI in healthcare. However, the predominantly positive framing I found in Danish coverage appears to echo a broader trend in Western media.

A 2021 peer-reviewed study titled The AI doctor will see you now: assessing the framing of AI in news coverage[38] analyzed how three major outlets — the British newspapers *The Guardian* and *The Daily Telegraph*, along with the American *Wall Street Journal* — reported on artificial intelligence in healthcare from 1980 to 2019. Based on an analysis of 365 articles, the researchers concluded that these publications largely portrayed AI in healthcare in a positive light.

[38] Mercedes Bunz and Marco Braghieri. 2021. “The AI doctor will see you now: assessing the framing of AI in news coverage.” *Springer Nature Link*

Many of the articles described AI as more precise, faster, and more effective than human experts. According to the authors, this has contributed to an *outperformance* narrative, where AI is not merely seen as a supporting tool, but as a technology that surpasses human judgment.

The study's authors argue that the media's framing of AI in healthcare is becoming increasingly important as the technology is deployed more widely. They stress the need for a well-informed public discussion on the benefits, limitations, and ethical challenges of AI in healthcare. A conversation that media outlets are well-positioned to facilitate.

However, they caution that consistently portraying AI systems as superior to human expertise risks placing them *above critique and concern* and may suppress important ethical debates. As they write, such uncritical framing *may be covering up potential ethical issues* and could ultimately hinder a more nuanced, democratic discourse on the future of healthcare.

Positive framing - a persistent pattern

The study is from 2021 and covers media reporting that stretches back several decades. That is a long period, not least when dealing with AI, which is currently developing at great speed. In a rapidly evolving field like artificial intelligence it may seem outdated to rely on data that doesn't capture the most recent developments. But more recent studies specifically examining media coverage of AI in healthcare are hard to come by.

However, according to one leading expert on media and AI in general, the overall media narrative about AI as such has not changed significantly in recent years. As Rasmus Kleis Nielsen writes in a 2024 overview article for the Reuters Institute:

"Several pieces of research on media coverage of AI suggest that overall, news coverage (...) comes off as industry-led and generally positive, even uncritical."

Rasmus Kleis Nielsen served as Director of the Reuters Institute for the Study of Journalism at Oxford from 2018 to 2024 and is currently a professor at the Department of Communication at the University of Copenhagen. He remains affiliated with Reuters as a Senior Research Associate.

When I spoke to him in a follow-up interview, he emphasized that high-quality, independent technology journalism certainly exists. But he also pointed to a structural trend in how AI is covered in the media:

"Much of the volume in the coverage tends to be the uncritical, source-driven kind," says Rasmus Kleis Nielsen.

It first became clear in 2018, when Reuters examined how UK media cover AI.[39] The study found that 60 percent of news articles across outlets were based on industry products,

[39] J. Scott Brennen and Rasmus Kleis Nielsen, 2018, "An Industry-Led Debate: How UK Media Cover Artificial Intelligence." *Reuters Institute*

initiatives, or announcements. Twice as many sources quoted in the articles came from industry as from academia or other sectors. In a qualitative analysis[40] , Reuters suggested that AI coverage in the UK tended to *construct the expectation of a pseudo-artificial general intelligence: a collective of technologies capable of solving nearly any problem.*

Rasmus Kleis Nielsen concludes that the majority of articles on AI were - and largely still are - relatively uncritical and source-driven. They are often single-source stories in which journalists relay information from individuals with a vested interest in AI, typically highlighting positive results and the potential of the new AI solution they are developing.

“A lot of it may very well be true. That artificial intelligence can do a lot of things. But it’s not always true that it happens without major side effects and complications. And often, things haven’t necessarily gone quite as well as those who are eager to promote their work would like to suggest,” says Rasmus Kleis Nielsen.

A 2022 research report[41] on media coverage of AI in Canada reveals a similar trend. The Canadian researchers wrote that *tech news tends to be techno-optimistic*. Overall, they found that *very few critical voices are heard in legacy media in Canada*.

On that basis, Rasmus Kleis Nielsen is not surprised that my analysis of Danish-language mainstream media coverage of AI in healthcare shows that the vast majority of articles take an overwhelmingly positive angle:

“That’s very much in line with the findings research has made in other countries,” says the former director of the Reuters Institute for the Study of Journalism.

However, he notes that the proportion of industry sources with a direct financial stake in promoting a positive AI narrative is lower in the Danish healthcare articles than in AI coverage more broadly. Instead, doctors and researchers are more frequently represented.

“That part looks generally better, in the sense that there seems to be quite a significant focus on the medical perspective,”he says.

Still, doctors and researchers working with AI may also have an interest in generating enthusiasm for the field. Regardless, Kleis Nielsen calls for more nuanced reporting on AI in healthcare—reporting that also addresses topics like algorithmic bias, accountability when AI makes mistakes, the right to privacy, surveillance, and the opacity of so-called black box systems:

“I know it’s easy to be generous with other people’s time and effort, and that there are many topics where researchers will say to journalists: you should spend more time and contact more sources. But coverage would be more nuanced if journalists also spoke with, for example, civil society organizations or rights-based organizations that have a lot of insight and knowledge in this area. In addition, users and recipients of AI would also be able to contribute important knowledge”, says Rasmus Kleis Nielsen.

[40] J. Scott Brennen et al. 2020. “What to expect when you’re expecting robots: Futures, expectations, and pseudo-artificial general intelligence in UK news” *Sage Journals*
[41] Guillaume Dandurand et al. 2020. “Training the News: Coverage of Canada's AI Hype Cycle (2012–2021)”

In healthcare, the recipients of AI are typically the patients. And the fact that only one of the 53 sources across 24 articles in Danish-language daily newspapers is a patient has not gone unnoticed by the director of the interest group *Danske Patienter*. *(Danish Patients)* While not surprised, he is concerned:

“That confirms our own experience, and I think it is concerning,” says Morten Freil.

He sees enormous potential in deploying AI in healthcare, for example improved and faster diagnostics and decision support in relation to treatment choices, but he calls for a public debate that also acknowledges the dilemmas:

“Coverage of the great AI potential should go hand in hand with highlighting the concerns, wishes, and needs of citizens and patients, as well as fair coverage of the dilemmas raised by AI. This can provide important input for AI developers and politicians, and is also crucial to maintaining public trust,” says Morten Freil, whom I will return to in the next chapter, where he offers specific suggestions on how media coverage of AI can be improved.

AI-panic-journalism

While research shows that the dominant narrative in media coverage of AI remains overwhelmingly positive, a smaller, emerging trend is worth noting.

A 2024 research article[42] published on Nature.com presents preliminary findings from a study of UK media headlines about AI technologies. The authors conclude that “media representations often swing between extremes of promising potential and serious impending dangers, including at times references to the outright annihilation of the planet.” In other words, articles with catastrophic framings are beginning to appear alongside the predominantly positive ones. According to the researchers, this could signal a more polarized media landscape, which might itself pose risks: “The at times bipolar representation of AI as both hero and demon may lead to unreasonable or inaccurate representations of the capabilities and functions of AI at present in society,” write Jasper Roe and Mike Perkins.

Rasmus Kleis Nielsen, the former director of the Reuters Institute, highlights the work of researcher Nirit Weis-Blatt, who sees the rise in AI doomsday scenarios as a reflection of the old journalistic maxim: “if it bleeds, it leads.” In short, stories about disasters, dystopias, and scandals generate clicks. Weis-Blatt calls this phenomenon “AI panic”-journalism-coverage that, like overly celebratory reporting, lacks essential nuance, according to Nielsen.

In his words, it is “journalism that tends to foreground claims about the possible extinction-risk that could arise from AI in the future while, effectively, diverting attention away from current, real-world problems ranging from discrimination and inequality to the environmental impact of energy- and water-hungry technologies.”

[42] Jasper Roe and Mike Perkins. 2024. “‘What they’re not telling you about ChatGPT’: exploring the discourse of AI in UK news media headlines.” *Nature*

As Nielsen points out, this kind of science-fiction-style catastrophe framing can obscure the more grounded and pressing debate about real-world issues, such as algorithmic bias, unequal access, or accountability in healthcare systems.

Still, despite the emergence of this countertrend, the overarching tone in AI coverage remains strongly positive. In the "Handbook of Critical Studies of Artificial Intelligence"[43] (November 2023), chapter 24 reviews a range of published studies on AI in the news, primarily across Western media. The conclusion: AI is receiving *largely positive evaluations and economic framing*, and the volume of coverage is rapidly increasing.

Drawing on these and similar findings, Rasmus Kleis Nielsen concludes that: "The coverage tends to be led by industry sources, and often takes claims about what the technology can and can't do, and might be able to do in the future, at face value in ways that contribute to the hype cycle." He adds a more concise assessment: "News coverage, often uncritical, helps build up the AI hype." And as he notes, this kind of coverage could ultimately come at a high cost.

Hype cycle and inflated expectations

The concept of the hype cycle, as referenced by Rasmus Kleis Nielsen, was developed in the 1990s by the American research and advisory firm Gartner to illustrate what they identified as a typical maturity path of emerging technologies, through five key phases:

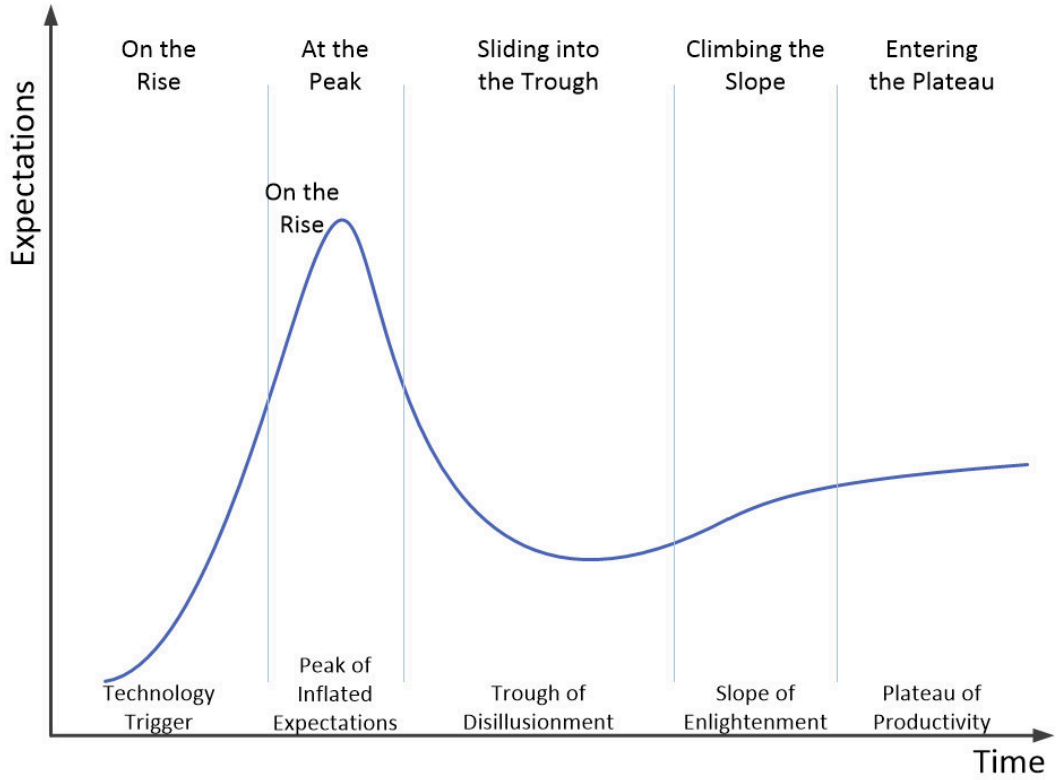

[43] Simon Lindgren and Edward Elgar. 2023 "Handbook of Critical Studies of Artificial Intelligence." *Edward Elgar Publishing*

1. Technology Trigger: A breakthrough innovation sparks early interest. Initial demos and media buzz create excitement, though functioning products are rare and success remains uncertain.

2. Peak of Inflated Expectations: Hype grows rapidly, driven by early success stories—often alongside many failures.

3. Trough of Disillusionment: Expectations drop as the technology fails to deliver on inflated promises.

4. Slope of Enlightenment: More practical use cases emerge, and understanding of the technology deepens. Expectations rise again—but this time at a more realistic level, aligned with the technology's actual potential.

5. Plateau of Productivity: The technology proves its value. It sees broad adoption, and expectations now match what it can deliver in real-world contexts.

According to Rasmus Kleis Nielsen, today's journalistic coverage of AI plays a role in boosting the hype to an unrealistically high level. If the hype cycle graph indeed reflects a typical trajectory for new technologies, and if my own, admittedly limited, analysis of Danish media's overwhelmingly enthusiastic AI coverage is accurate, then we are currently positioned high on the curve, in the phase known as the *Peak of Inflated Expectations*.

"If overwhelmingly positive media coverage inflates expectations to an unrealistic level, with anticipation of imminent technological breakthroughs, it attracts massive investments," says Rasmus Kleis Nielsen.

But when expectations exceed what the technology can actually deliver, the result is often an *AI winter* - a period of disappointment, in which AI systems fail to meet the hype. As excitement deflates and expectations slowly adjust to reality, investment value may drop drastically, with investors facing substantial losses. When it is public institutions investing, it is taxpayers' money at stake.

"We've already lived through several AI winters, and we've seen the same dynamic with other large-scale public IT projects," says Rasmus Kleis Nielsen. That's one reason, he argues, for a more nuanced and realistic media coverage. A coverage that doesn't overinflate the prospects of AI.

Thomas Ploug, professor of data and AI ethics at Aalborg University and author of *Medical Ethics*, also believes that AI in healthcare will not only bring about both groundbreaking opportunities, but also inevitable failures -hallmarks of what Gartner defines as the *Trough of Disillusionment* (phase three of the hype cycle). As mentioned earlier in this report, Ploug expects that AI will generate scandals in the healthcare sector, akin to those we've seen in other areas, such as the crisis-plagued Danish real estate tax assessment system:

"Mistakes are bound to happen. And we'll see AI implemented in areas where it should never have been introduced," says Thomas Ploug.

Considering the research findings and media criticism discussed throughout this report, it does not seem unreasonable to conclude that AI in healthcare, too, calls for a more nuanced and layered form of journalism than what we currently see. When 22 out of 26 articles on AI and healthcare over a course of 15 months are overwhelmingly celebratory, when source diversity is relatively narrow, when patients rarely appear as voices in the coverage, and when debate and criticism are relegated to opinion pages and letters to the editor—then surely, there is room for stronger journalism. Journalism that fosters public dialogue, invites engagement, and makes space for nuance, critique, and realism—without losing sight of the many promising possibilities.

How can that be done? That’s the subject of the next chapter.

How the media can do better

In this chapter, I explore whether the core ideas and principles of constructive journalism can serve as inspiration to improve media coverage of AI in healthcare. And if so, how?

To start with what may seem like an obvious conclusion (at least to a fellow who has just spent 10 months at the Constructive Institute), my answer is: Yes, constructive journalism can be a lever for better AI journalism. With the addition: Especially if one keeps in mind the sometimes slightly overlooked credo within constructive journalism: 'Remember to investigate solutions with the same critical rigor as when uncovering the problems.'

The three main pillars of constructive journalism are: *focus on solutions, cover nuances,* and *promote democratic conversation.*

Let us look at how each of these pillars can be activated and support better coverage of AI in healthcare.

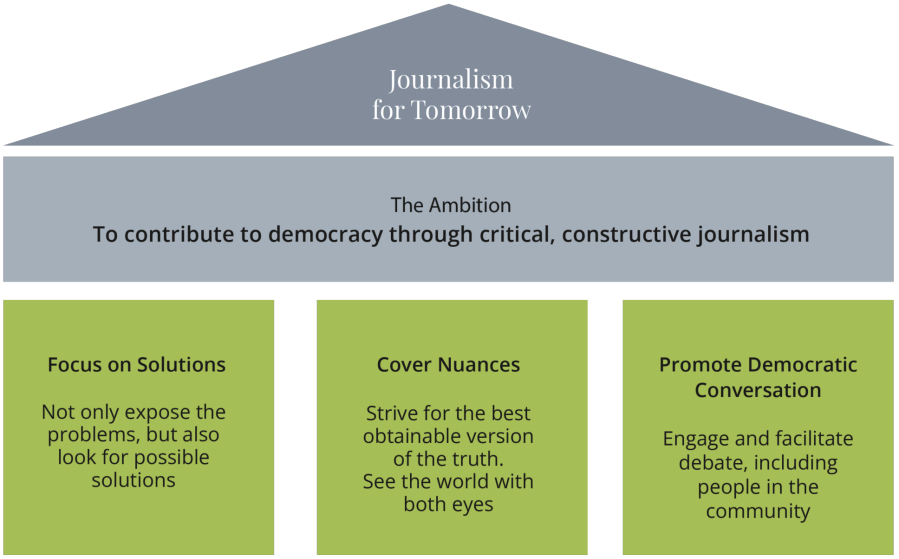

The three pillars of constructive journalism. Source: Constructive Institute website

Focus on Solutions - but don't forget the critical questions!

Constructive journalism, as defined by the Constructive Institute in Aarhus[44], is a "response to increasing sensationalism and negativity bias of the news media today." It focuses on solutions to problems, but "is rigorous and critical in its approach, both when it comes to reporting on problems and progress."

In other words, constructive journalism aims to counterbalance many media outlets' excessive focus on negative news by not just pointing to problems but also by investigating

[44] *Constructive Institute homepage: https://constructiveinstitute.org/why/*

solutions. At the same time, solutions and progress should not be accepted or conveyed without thorough, critical examination. As the Institute's website puts it: *"*[Constructive Journalism] doesn't engage in 'feel good' stories or look at the world through rose-tinted glasses. Instead, the goal is to vigorously report on possible solutions to important social issues."

This is an area where mainstream coverage of AI in healthcare in Denmark could improve. On the subject of AI, they seem to have a preference for rose-tinted glasses. Of the 26 articles examined in Danish mainstream daily newspapers, 22 had a one-sided, positive framing, largely focusing on how specific AI projects could solve a myriad of problems and challenges for patients and healthcare professionals in the Danish system. So the issue here is not a lack of focus on solutions. The issue is that very few - and in many cases no - critical questions are asked about the solutions presented.

In their publication Handbook for Constructive Journalism[45], Kristina Lund Jørgensen and Jakob Risbro describe the constructive method this way: *We identify a problem with known consequences. When we decide to focus on a possible solution, we approach it with the same sharp, critical attitude we would adopt when exploring a societal problem.* In the case of covering AI in healthcare, such an approach would strengthen the credibility of the news article and test the durability of the AI solution it covers in a trustworthy manner.

Put the solution through the test

Journalism about emerging AI in healthcare almost by definition focuses on opportunities, hopes, and solutions, especially in the phase we're currently in, where promising new AI systems appear nearly every week or month, many of which are not yet fully implemented in practice.

But even if these new AI systems haven't yet been deployed - or have only been in use briefly - journalists still have an opportunity to ask critical questions. These can help media users assess the new technology and form hopes (or skepticism) on a more realistic and credible basis.

Here's a battery of specific questions journalists might consider asking when presented with a new breakthrough and an AI system designed to improve tasks in the healthcare sector. This list is meant for inspiration and questions should, of course, be chosen, shaped, supplemented or expanded depending on the specific context:

- What exactly can the system do—and what can it not do?
- How does it differ from existing solutions?
- What data is the system trained on—and from which populations? *(= check for bias)*
- Has patient testing or user involvement been part of the development phase? *(= check whether user needs have been considered and whether the AI was adjusted*

[45] Kristina Lund Jørgensen and Jakob Risbro. 2021. "Handbook for Constructive Journalism." *Constructive Institute and IMS Media Learning Hub*

based on real-life testing)

- Has the system been validated in clinical practice—and in which settings? *(= has it been proven to work outside the lab?)*

- Has it been adapted to the Danish context in which it will be used? *(Foreign AI systems often require adjustments to local legislation, software, workflows in the Danish healthcare regions, patient characteristics, etc.)*

- Is the AI system scalable for nationwide implementation?

- What documented test results are available?

- How is the risk of bias handled—e.g. in relation to age, gender, ethnicity, or socioeconomic status?

- Have any independent third-party evaluations been conducted? *(= not just evaluation by those with a vested interest in the system)*

- How is patient consent handled when AI is used?

- What data sources does the system use—and how is personal data protection ensured?

- Are the algorithm's decisions explainable to healthcare professionals and patients?

- Is there a public and transparent description of how the system works - or is it a black box?

- How will limitations and error margins be communicated to users?

- Who pays if the AI system does not perform as expected?

- Who is held accountable if the AI makes faulty decisions?

Seek a second opinion

Another obvious way to assess an AI solution is by seeking a second opinion from an independent source. Many AI solutions are presented by developers or healthcare organizations with economic, political, or branding interests. A second opinion can help ensure that, as a journalist, you don't end up uncritically amplifying hype or marketing spin.

Independent experts may point out methodological flaws, bias, untested assumptions, or unintended consequences that might otherwise go unnoticed and are unlikely to be mentioned in press releases or product presentations.

Bringing in an independent expert for a second opinion enhances the article's credibility and builds reader trust that the journalism is thorough, nuanced, and free from spin. This is particularly important in AI healthcare journalism, where readers often have limited technical knowledge. Moreover, a second opinion can sometimes reveal that multiple professional perspectives exist, underscoring that AI in healthcare is not an exact science. It contributes to a more realistic understanding of the world. The appropriate type of source for a second opinion depends on the specific healthcare AI system discussed in the article. Depending on the topic, consider sources such as:

- Clinical researchers unaffiliated with the project
- Academics specializing in medical ethics or clinical AI
- Data protection and patient safety specialists
- Members of ethics councils
- Organizations representing patient groups.

There are very few if any articles in my review of Danish daily newspapers that simultaneously explore both potential and dilemmas, cover nuances, and include sources such as patients, physicians, industry professionals, and ethicists. But it is possible! Jyllands-Posten demonstrates this with a nuanced story on the front page and in a full spread of its Sunday edition on June 1, 2025. The article was published after the deadline for the formal study of this report, but it deserves mention here as a good, constructive example of how it can be done.

SØNDAG DEN 1. JUNI 2025 | JP.DK | KR. 59,00

MORGENAVISEN

Jyllands-Posten

SØNDAG

Israel er i stilhed ved at forberede en **annektering** af Vestbredden.
INDBLIK, side 6-8

Caroline Nieling ligner ikke en alkoholiker. Men i 10 år drak hun for at komme igennem dagen.
Side 10

Hver 10. danske læge bruger AI

Kunstig intelligens (AI) har gjort sit indtog i danske læge-huse. Det rejser en række alvorlige dilemmaer.
Side 6-7

Læge Christian Heilmann bruger AI i sin praksis i Lynge. Computeren optager hans samtaler med patienterne. Samtalerne transkriberes, så han kan bruge notaterne i sin journalføring.
Foto: Stine Bidstrup

Jyllands-Posten june 1st 2025: One out of every 10 Danish doctors uses AI

The article reports that more than one in ten general practitioners in Denmark have now purchased an AI assistant – an artificial intelligence system that listens in during consultations between doctor and patient. After the consultation, the AI produces a full transcript of the conversation and generates a draft journal entry that the doctor can review and insert into the patient's medical record.

This potentially saves the doctor a great deal of time – time that could benefit the patient. However, it also raises several important ethical concerns. What happens, for example, if the system makes mistakes and inserts inaccurate or incorrect information into the record without the doctor noticing? Just to name one issue.

The article includes voices from a general practitioner, a professor in data and AI ethics (Thomas Ploug, who is also interviewed in this report), a sales director from the AI assistant company, the chairman of the Danish Association of General Practitioners, and a patient. This provides a rich and nuanced perspective in an article that presents both the AI solution and its associated dilemmas, allowing readers to make up their own minds. Well done!

Do follow ups – it pays off!

Once a new AI solution has been implemented in a hospital and the initial hype surrounding its potential has faded, it's time for follow-up reporting. There are important and insightful stories to be found here.

The Handbook for Constructive Journalism[46] highlights this as a key principle: "Once we identify any replicable, scalable solution(s), we present our findings to those responsible for solving the problem. **After a while, we follow up to see if the issue has been resolved**." (my emphasis added)

Rasmus Kleis Nielsen, professor at the Department of Communication at the University of Copenhagen and former director of the Reuters Institute for the Study of Journalism, agrees:

"For me, follow-up is one of the most essential parts of journalistic coverage of AI. It's where you assess whether the promises actually hold water."

Much AI journalism, he argues, is built around optimistic promises made at the launch of a new system. But it is equally important to revisit those promises once the system has been put to work in the real world:

"That's when you're dealing with things that have actually happened and can be independently verified. At that point, you're not relying solely on those with a vested interest."

In my analysis of 26 articles from Danish language daily newspapers covering AI in healthcare, I found just one example of follow-up journalism.

[46] Kristina Lund Jørgensen and Jakob Risbro. 2021. "Handbook for Constructive Journalism." *Constructive Institute and IMS Media Learning Hub*

Gribskov Kommune ser et stort potentiale i kunstig intelligens. Foto: Thomas Arnbo

Sådan skal AI spare ældreplejen for millioner

Gribskov: Allerede i 2024 er det planen, at kunstig intelligens skal reducere udgifterne med fire procent.

Skulle spare kommunen for millioner - stort AI-projekt må droppes

Gribskov: Leverandøren af et nyt hjælperedskab med kunstig intelligens kunne ikke levere. Nu har kommunen trukket stikket.

Efter fire forsinkelser har kommunen trukket stikket på projektet. Foto: Nicklas Thorup Clausen

Januar 2024; AI will save Gribskov millions af kroner

Juli 2024: The project did not work and was dropped.

The article in question, published by Frederiksborg Amts Avis in July 2024, was titled

- *Skulle spare kommunen for millioner – stort AI-projekt må droppes* (*Supposed to save the municipality millions – major AI project scrapped*).

It followed up on a story from January the same year

- *Sådan skal AI spare ældreplejen for millioner* (*How AI is set to save millions in elderly care*).

The original article reported that Gribskov Municipality was in the process of implementing an AI-assisted shift planning system, which, according to the municipality, would optimize staff schedules in elderly care and save six million Danish kroner that year alone. Equivalent of roughly four percent of the overall budget for home care services. However, the follow-up piece revealed that the system never worked as intended, and the municipality eventually terminated its contract with the vendor.

It could have been just as good a story if the follow-up had shown that the system exceeded expectations and delivered twice the savings projected. The point is: when AI meets reality, its true potential becomes visible. Follow-up journalism helps test whether the technology is actually making a difference in practice, and contributes to a more accurate picture of both opportunities and limitations. Follow up articles provide a window into how AI is reshaping workflows, responsibilities, trust, and decision-making processes. They also open the door to constructive questions about how AI systems can be improved moving forward.

As a checklist, journalists might consider approaching follow up journalism with at least six key questions in their mind:

- What was promised?
- What was delivered?
- What does the data say about the results?

- What do the users say?
- What have we learned?
- What happens next?

Even now, some follow-up reporting can already be done. And as AI systems continue to roll out across the healthcare sector in the months and years to come, this should become an even bigger priority. There's a goldmine of constructive, citizen-centered stories waiting to be told. Stories that go beyond technological “breakthroughs” and instead paint a documented picture of how AI solutions are actually affecting healthcare on the ground.

Cover nuances - listen to different voices

The second pillar of constructive journalism is “Cover Nuances.” In essence, this principle is about avoiding oversimplification and “seeing the world with both eyes.” The goal is to offer audiences a balanced and more accurate representation of reality.

Previous sections in this chapter have outlined several journalistic practices that contribute to more nuanced coverage of AI in healthcare: asking critical questions about proposed solutions, seeking second opinions, and producing follow-up stories. Here, I focus on another vital strategy: expanding the range of sources to include more diverse voices, particularly those of users and patients.

Let’s start with the patients. Among the 53 sources cited in the 26 articles I analyzed on AI and healthcare in Danish newspapers, only **one** was a patient. Additionally, representatives from the patient organization Danish Cancer Society (Kræftens Bekæmpelse) appeared in three articles. In the remaining articles, the patient perspective is notably absent.

As Rasmus Kleis Nielsen, former Director of the Reuters Institute for the Study of Journalism, dryly observes:

“In SO many stories about AI in healthcare, it would be refreshing to see a few more citizen perspectives.”

In a 2024 article on AI ethics in the medical periodical Ugeskrift for Læger[47] three experts affiliated with Odense University Hospital underscore the importance of involving patients in the development of healthcare-related AI. Anne Gerdes (CAI-X), Iben Fasterholdt (Center for Innovative Medical Technology), and Benjamin S.B. Rasmussen (CAI-X and Department of Radiology) write about the patients:

“Their perspectives can help identify potential ethical challenges and contribute to a fair and responsible development and implementation of AI in the healthcare system…Patient involvement can help ensure that patient needs are better met and that public trust in research increases.”

[47] Anne Gerdes et al. 2024. “AI-etik i sundhedsvæsenet.” *Ugeskrift for læger*

This applies not only to the development of AI, but also to journalism about it. That view is shared - perhaps unsurprisingly - by Morten Freil, Director of the Danish Patients' Association. He is not shocked to learn that only one patient and three patient organization representatives appear among the 53 sources in the articles I analyzed:

"Those are some striking numbers, but honestly, they don't surprise me. That's also our experience. I find it somewhat concerning," says Morten Freil.

While Freil sees tremendous potential in the widespread use of AI in healthcare and strongly supports continued research and development, he also calls for more balanced media coverage that includes not just optimistic visions but also patients' needs, doubts, and concerns:

"There are a number of dilemmas in this development that need to be taken into account—and we believe they receive far too little attention in public debate."

Among the concerns he mentions are: the risk of bias in AI systems; the fear that human contact between patient and provider could be reduced if too many decisions are delegated to AI, and the question of who holds responsibility if an AI-supported decision ends up harming a patient.

"It's crucial for patient trust that their concerns are voiced and addressed. That's essential both during the development of AI and in the public debate that the media help facilitate" says Morten Freil.

Journalists covering AI and healthcare would benefit from establishing a broader source network that includes patient voices and organizations representing them. These might include the Danish Cancer Society, Ældre Sagen (DaneAge Association), the OCD Association, and the umbrella group Danish Patients.

Beyond the medical, political, and technical sources typically consulted, journalists should also consider voices from civil society and human rights organizations. Especially when topics like privacy, surveillance, and algorithmic bias are on the table. As former Reuters-director Rasmus Kleis Nielsen noted in the previous chapter:

"Coverage would be more nuanced if journalists also spoke with, for example, civil society organizations or rights-based organizations."

Of course, it's not always possible to consult every type of source in every article. But reporters should at least take a moment to consciously consider who is best placed to bring nuance and depth, given the time and context.

A practical tool might be to ask: Which of these four types of sources could be relevant to include in the story about AI in healthcare, I am about to do: A healthcare source, a patient representative, an ethical or legal expert, a tech source.

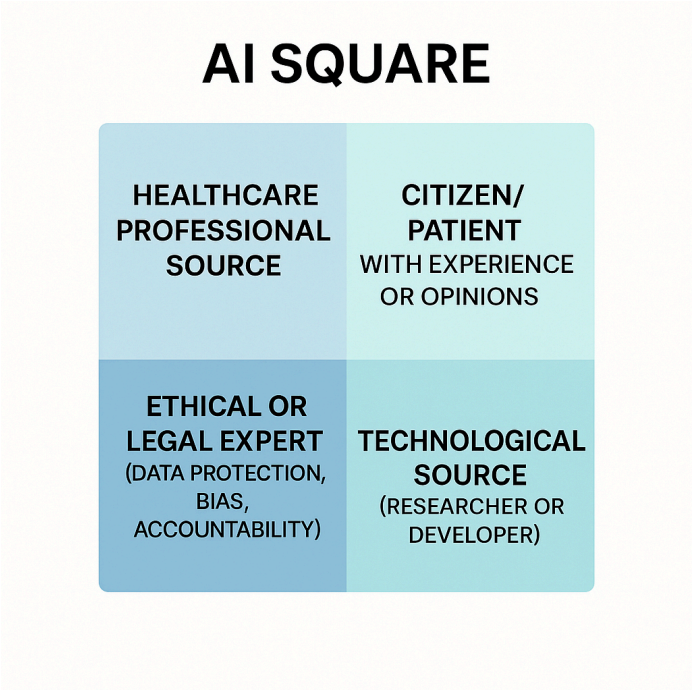

Here visualized in my homemade memory model, “The AI Square” :-)

Promote democratic conversation

The third pillar of constructive journalism is to *promote democratic conversation.* To engage the public and facilitate debate by including people from the community. The goals include uncovering citizens' actual concerns, hopes, and needs, identifying pressing problems, and inviting public dialogue around potential solutions. According to the Danish-Norwegian white paper on constructive journalism[48], published by the University of Southern Denmark in October 2024, the aim is not to push one solution as the right one, but rather “to facilitate constructive debate that seeks to answer the ‘what now and how’ based on journalistic research.”

With this mindset, it makes obvious sense to draw on the tools of the third pillar when reporting on a field as future-oriented and dilemma-filled as artificial intelligence in healthcare.

Uncover Public Concerns, Needs, and Hopes

Dialogue-based journalism is a promising approach to uncover what citizens think, engage them in the early stages of the journalistic process, and encourage democratic debate. It’s a model where the core is to listen actively and involve users throughout the process — from story idea to feedback. In 2022, the journalism program at DMJX in Aarhus published an extensive guide to inspiration on dialogue-based journalism[49], defining it as journalism where *“the common denominator is listening to users and including them from idea phase to comment thread.”*

[48] Orla Borg et al. 2024. “Konstruktiv journalistik. En dansk-norsk hvidbog”. *Syddansk Universitet*
[49] DMJX et al. 2022. “Dialogue-based Journalism – An Inspirational Guide.”

Below are some concrete examples that illustrate this way of thinking. I have not tailored or reworked them to specifically fit coverage of AI in healthcare, so they should be seen more as inspiration that might spark further concept development on how the principles can be applied in concrete AI coverage - especially when the media aim to engage users, understand citizens' perspectives on AI in healthcare, and report more thoroughly on the subject.

TV 2 Kosmopol – "Ask Us!"

"Ask Us!" is a concept developed by TV 2 Kosmopol, which covers two Danish regions and 34 municipalities in the greater Copenhagen area. The concept is structured in four steps:

Illustration fra tv2kosmopol.dk: 1 Submit your question, 2. Vote for your favorite, 3 We investigate, 4. We share the answer

1. **Submit your question.** On Kosmopol's website, users are encouraged: "Ask us! What are you wondering about in the Capital Region?" They are directed to a submission box where they can write in their questions for the newsroom to investigate.

2. **Vote for your favorite.** Each week, the editorial team selects three questions. Viewers vote online for the one they're most curious about.

3. **We investigate.** Journalists research and report on the winning question.

4. **We share the answer.** TV 2 Kosmopol publishes a segment or article answering the question.

Questions range from the hyperlocal ("How often are grocery store baskets cleaned – if ever?") to the civic and structural ("Why don't we cover ugly motorways with solar-powered noise barriers?").

During a presentation at Constructive Institute in March 2025, Nanna Holst — editor for dialogue and engagement at TV 2 Kosmopol — noted that the format provides a unique window into what the audience genuinely cares about. It gives journalists access to story ideas rooted in public interest all the while engaging the users in dialog and journalism. When the newsroom listens to its users instead of deciding unilaterally what matters, engagement rises significantly, Nanna Holst says.

Nanna Holst, herself a Constructive Institute fellow, also authored the report Better listening[50], which offers extensive practical advice on how media outlets can improve their ability to listen and uncover public needs.

A media that wishes to cover AI in healthcare from a patient/citizen point of view, could, do something similar and find a way to get access to the patients' wonder and questions and let it guide journalistic projects on the subject.

Zetland – Ask the members

The Danish digital media outlet Zetland refers to its paying users not as subscribers but as members. It is a deliberate strategy meant to signal more than just a transactional relationship. It sends a message that readers are seen as part of a shared project, perhaps even a movement for quality journalism.

Members are frequently invited into the editorial process. On the Zetland site, the newsroom regularly posts questions such as:

- "What do you want our journalists to investigate?"
- "Can you help research a story?"
- "Do you know relevant sources we should interview?"

Members are also encouraged to comment on stories on the website, provide constructive feedback, engage in discussions with journalists and one another, and suggest follow-up angles.

In a video interview with DMJX[51] in 2022, then editor-in-chief Ida Ebbensgaard explained that member involvement improves the quality of journalism and generates unexpected and valuable story ideas:

"They very often know more than us on a subject. It is extremely rewarding to tap into this wide range of knowledge. It gives wonderful journalism."

Perhaps other media could reach out to its users in a similar way in order to improve reporting on AI in health.

Open Editorial Meetings

Some news outlets - from the local newspaper *Midtjyllands Avis* to the policy-focused *Mandag Morgen* - have experimented with open editorial meetings.

At *Midtjyllands Avis*, readers were invited to sit in on daily meetings, hear journalists' story considerations, and suggest topics for coverage. *Mandag Morgen* created thematic

[50] Nanna Holst. 2023. "Better listening. Shortcut to a more listening mindset." *Constructive Institute*
[51] DMJX et al. 2022. "Dialogue-based Journalism – An Inspirational Guide."

hour-long sessions open to subscribers, where they could engage in conversations around questions such as:

- “Who should own and manage our data?”
- “What if low growth is the new normal?”
- “How can we get more marginalized youth into work and education?”

These sessions were held early in the editorial process so that user contributions could meaningfully shape the final reporting. Journalists facilitated and listened rather than leading the conversation.

This model could also be adapted into early-phase idea sessions with invited stakeholders, experts, or interest groups — for example, in the context of a planned reporting series on AI in healthcare.

Citizen Hearings

In February 2024, TrygFonden hosted a public hearing in collaboration with the Ministry of Digitalization. The goal was to spark a democratic conversation around artificial intelligence and provide input to policymakers about public concerns and expectations.

A group of 350 citizens was selected to represent the Danish population. They received background material, participated in group discussions, and had the chance to question experts. The event was covered by *Politiken* and summarized in a public report.

While this kind of large-scale effort requires substantial funding and organization, the underlying concept of evolving citizens can inspire to do similar projects on a smaller scale, thus creating involvement and journalism at the same time: Media outlets could host community conversations, cover the events before and after, and gather input for future reporting on AI and health.

Some of the suggestions in this chapter, such as checklists and simple sourcing routines, are easily implemented. Others, like public engagement formats, may be more time- or resource-intensive. Think of this chapter as an inspirational toolbox. Not every tool will fit every newsroom or context, but there’s value in selecting what works based on the avaiilable time, ambition, and opportunity.

Sometimes, a quick day-to-day news piece can be strengthened by just one extra source. At other times, larger editorial projects can benefit from more elaborate forms of public engagement.

Final words

All in all there are many ways journalistically to involve and engage the public in the discourse, to look for nuances, and to critically examine constructive solutions to societal challenges. Ultimately it begins with applying a mindset that it is important. And with a willingness to cover the AI revolution in health care in a balanced way in order to enlighten the public debate and give decisionmakers a sense of the hopes, worries and prioritizations of the people. At the end of the day we need a better understanding of both the huge potentials and the big dilemmas that accompany AI in healthcare in order for society to guide and shape the future implementation. It is, quite literally, a question of life and death.

Appendix

Analysis of Danish-language daily newspapers' coverage of AI in healthcare in the period from January 1, 2024, to April 1, 2025.

Total number of articles	26
Positively framed articles	22
Neutrally framed articles	2
Critically framed articles	2
Sources used in the articles	53
Experts (researchers, doctors, AI-specialists)	24
Industry representatives	8
Politicians	8
Civil servants	6
NGO/Interest groups	5
Patients	1

Overview of headlines, media, and publication dates of the analyzed articles
(The headlines are quoted as they appear in the Infomedia search. They may differ slightly from the printed versions or the media's online editions)

Positively framed articles:
Hans metode kan fastslå risikoen for kræft (His method can determine the risk of cancer), Berlingske 18/1 2024

Sådan skal AI spare ældreplejen for millioner (*This is how AI will save eldercare millions*), Frederiksborg Amts Avis 29/1 2024

Første menneske har fået Musk-chip i hjernen (First human receives Musk brain chip), Børsen 5/2 2024

Computer opdager kræft, hjertestop og brækkede ben. AI stormer frem på hospitaler (*Computer detects cancer, cardiac arrest, and broken bones. AI surges in hospitals*), Berlingske 22/2 2024

Manglende lovgivning står i vejen for brug af kunstig intelligens i ældreplejen (*Lack of legislation hinders the use of artificial intelligence in eldercare*), Berlingske 3/3 2024

Kunstig intelligens bliver et tungt kræftvåben (*Artificial intelligence becomes a powerful cancer weapon*), Flensborg Avis 11/4 2024

Kunstig intelligens kan vise en vej ud af kræftkrisen (Artificial intelligence may show a way out of the cancer crisis), Berlingske 19/4 2024)

Pensionschef: Patienter kan få afgørende hånd af kunstig intelligens (Pension director: Patients may receive crucial help from artificial intelligence), Berlingske 29/4 2024

Robotterne rykker tættere på din hospitalsseng (Robots move closer to your hospital bed), Flensborg Avis 23/5 2024

Læger skal bruge kunstig intelligens til at afsløre frygtede kræftsygdomme (Doctors to use artificial intelligence to detect feared cancers), Berlingske 24/5 2024

Danmark skal gøre kræftkampen kunstigt intelligent (*Denmark must make the cancer fight artificially intelligent*), Aarhus Stiftstidende 1/6 2024)

Ny forskning med udstrakt hånd til multisyge ældre (New research lends a helping hand to elderly patients with multiple illnesses), Dagbladet Roskilde 19/8 2024

Kunstig intelligens kan forudsige, hvem der har høj risiko for at blive tvangsindlagt (Artificial intelligence can predict who is at high risk of involuntary hospitalization), Jyllands-Posten 28/11 2024

Kæmpe potentiale (Huge potential), Ekstra Bladet 2/11 2024

Nu kan kunstig intelligens forudsige tvangsindlæggelser (Now artificial intelligence can predict involuntary admissions), Kristeligt Dagblad 7/12 2024

En kunstig intelligens registrerer alt på plejehjemmet, (Artificial intelligence monitors everything in the nursing home), Jyllands-Posten 29/12 2024

Den nye hjælper hedder AI (The new helper is called AI), Folkebladet Lemvig 18/1 2025

Da en ung studerende døde af tarmkræft, fik Ismail Gögenur nok (When a young student died of colon cancer, Ismail Gögenur had had enough), Politiken 10/2 2025

Kunstig intelligens skal hjælpe mænd med hurtigere kræftsvar (Artificial intelligence to help men receive faster cancer diagnoses), Nordjyske Stiftstidende 4/2 2025

EU vil rejse 200 milliarder euro til investering i AI, (EU plans to raise 200 billion euros for investment in AI), Midtjyllands Avis 12/2 2025

Kunstig intelligens har lokaliseret molekyle, der kan knockoute PRRS (Artificial intelligence has identified a molecule that could knock out PRRS), Effektivt landbrug 20/2 2025

Danske AI-virksomheder møder muren hos jurister i den offentlige sektor , (Danish AI companies hit a legal wall in the public sector), Jyllands-Posten 28/3 2025

Critically framed articles

Skulle spare kommunen for millioner - stort AI-projekt må droppes, (Supposed to save the municipality millions – major AI project must be scrapped, Frederiksborg Amts Avis 22/7 2024

Skjult bag skærmen (Hidden behind the screen), Weekendavisen 6/9 2024

Neutrally framed articles

Forskere vil forebygge selvmord med AI (Researchers aim to prevent suicide with AI, Flensborg Avis 24/8 2024

SF og PFA i historisk alliance: AI skal gøre danskerne sundere (*SF and PFA in historic alliance: AI to improve*

Overview of the 53 sources used in the articles

Experts (researchers, doctors, AI specialists)

Søren Udby, Head of Center for Clinical Robotics (OUH + SDU)
Kjeld Møller Pedersen, Health Economist, University of Southern Denmark
Peter Børker Nielsen, Head of CAI-X, OUH's Center for Artificial Intelligence
Peter Børker Nielsen, Head of CAI-X, OUH's Center for Artificial Intelligence
Søren Dinesen Østergaard, Professor of Clinical Medicine, Aarhus University and Psychiatry in the Central Denmark Region
Søren Dinesen Østergaard, Professor of Clinical Medicine, Aarhus University and Psychiatry in the Central Denmark Region
Dorte Nielsen, Chief Nurse, Regional Hospital in Hjørring
Peter Buss Lasborg, Head, Department of Radiology, Aalborg University Hospital
Søren Pihlkjær Hjortshøj, Chief Medical Director, Aalborg University Hospital
Søren Pihlkjær Hjortshøj, Chief Medical Director, Aalborg University Hospital
Michael Borre, Senior Consultant, Professor, Department of Urology, Aalborg University Hospital
Ismail Gögenur, Professor and Surgeon, Zealand University Hospital, Køge
Benjamin Rasmussen, Clinical Research Lead, CAI-X, OUH's Center for Artificial Intelligence
Benjamin Rasmussen, Clinical Research Lead, CAI-X, OUH's Center for Artificial Intelligence
Nanna Winther Selmer, Nurse, PhD student
Josefine Freiberg, Medical Doctor and PhD, University of Copenhagen
Ditte Sloth Møller, Medical Physicist, Oncology Department, Aarhus University Hospital
Lise Bech Jellesmark Thorsen, Oncologist, Aarhus University Hospital
Carsten Utoft Niemann, Senior Consultant, Research Project Leader, Rigshospitalet
Ivan Branslund, Professor of Artificial Intelligence, University of Southern Denmark
Serkan Ayvaz, Associate Professor, Center for Industrial Software, SDU
Merete Nordentoft, Psychiatrist
Rasmus Birk, Psychologist and Researcher, Aalborg University
Michael Eriksen Benros, Professor, Research Director, Psychiatric Center Copenhagen

Industry representatives

Søren Gershøj, Sales Director, Omilon (provider of IT tools, especially to the healthcare sector)
André Rogarczewski, CEO, Netcompany
Mikkel Wad Thomsen, CEO, Teton-ai
Gitte Amstrup Sandlykke, CEO, Pondoo
Michael Guldager Hansen, Senior Sales Specialist, KMD
Elon Musk, CEO, Neuralink
Jesper Bjerre, COO, Danica
Ole Krogh Petersen, CEO, PFA

Politicians

Ursula von der Leyen, President, European Commission
Sisse Marie Welling, Chair, KL's Committee on Health and Elderly Affairs
Mette Kierkegaard, Minister for the Elderly
Gunnar Lisby Kjær, Chair, Social and Health Committee, Lemvig Municipality
Mads Duedahl, Chair, Region of North Jutland

Sophie Løhde, Minister of the Interior and Health
Lars Gaardhøj, Chair, Regional Council, Capital Region
Pia Olsen Dyhr, Chair, Socialist People's Party (SF)

Civil servants

Louise Halgaard Godtfredsen, Chief Consultant, South Denmark Innovation under Danish Regions
Per Nielsen, Head of Elderly Services, Lemvig Municipality
Kurt Espersen, Executive Director, Region of Southern Denmark
Kurt Espersen, Executive Director, Region of Southern Denmark
Mette Bierbaum, Head of Center for Health and Care, Gribskov Municipality
Miriam Toft, Director of Social, Health and Care, Gribskov Municipality

NGO/Interest groups

Janne Bisgaard, Senior Consultant, Project Director for Screening, Danish Cancer Society
Pernille Slebsager, Head of Department, Danish Cancer Society
Pernille Slebsager, Head of Department, Danish Cancer Society
Karin Zimmer, General Practitioners' Organization, PLO
Jeppe Kristen Toft, Director, Lifeline

Patients

Jonathan Løw, former psychiatric patient