

# Kunstig Intelligens og sundhed



# Tidsskrift for Forskning i Sygdom og Samfund

Nr. 44: Kunstig Intelligens og sundhed

© 2026 forfatterne.

## Redaktion:

Sara Marie Hebsgaard Offersen (ansvarshavende redaktør), Steno Diabetes Center Aarhus  
Mette Bech Risør (ansvarshavende redaktør & økonomiansvarlig), Institut for Folkesundhedsvidenskab,  
Center for Almen Medicin, Københavns Universitet og Institut for samfunnsmedisin,  
Allmenmedisinsk forskningsenhet, UiT Norges arktiske universitet.  
Torsten Risør (temareaktør), Afdeling for almen medicin og Forskningsenhet for almen praksis, Institut  
for Folkesundhedsvidenskab, Københavns Universitet og Fagenhet for primærmedisin, Institut for  
samfunnsmedisin, UiT Norges arktiske universitet.  
Marie Louise Tørring, Afdeling for Antropologi, Aarhus Universitet.  
Hanne Mogensen, Institut for Antropologi, Københavns Universitet  
Gitte Wind, Professionshøjskolen København, København  
Ann Dorrit Guassora, Forskningsenheten for Almen Praksis, Københavns Universitet  
Claus Bossen, Digital Design og Informationsvidenskab, Institut for Æstetik og Kommunikation, Aarhus  
Universitet, Danmark  
Camilla Merrild, Center for Almen Medicin, Aalborg Universitet  
Jeppe Oute, Institut for Mennesker og Teknologi, Roskilde Universitet  
Stinne Glasdam, Hälsa i ett tvärvitenskapligt perspektiv, Lund Universitet  
Martin Marchman Andersen, Statens Institut for Folkesundhed, Syddansk Universitet  
Ursula Marie Dyekjær (sekretær), Afdeling for Antropologi, Aarhus Universitet

## Gæstereaktør:

Claus Bossen, Informationsvidenskab, Aarhus Universitet  
Maja Hojer Bruun, Pædagogisk Antropologi, Aarhus Universitet  
Sune Hannibal Holm, Sektion for Forbrug, Bioetik og Regulering, Københavns Universitet

Tidsskriftet er udgivet med støtte fra: Det Frie Forskningsråd

Peer review foretages af et tværvidenskabeligt panel bestående af bl.a. læger, antropologer, filosoffer,  
historikere, psykologer, politologer og sociologer.

Proof, layout og prepress: Ursula Marie Dyekjær

Udgiver: Foreningen Medicinsk Antropologisk Forum. Institut for Kultur og Samfund, Antropologi,  
Aarhus Universitet, Moesgård, 8270 Højbjerg

Henvendelser: Tidsskrift for Forskning i Sygdom og Samfund.

Institut for Kultur og Samfund, Antropologi, Aarhus Universitet, Moesgård, 8270 Højbjerg.  
kontakt@sygdomsamfund.dk

ISSN (online): 1904-7975

## Formål:

*Tidsskrift for Forskning i Sygdom og Samfund er et tværfagligt tidsskrift, der tager udgangspunkt i medicinsk antropologi. Tidsskriftet har til formål at fremme og udvikle den forskning, der ligger i grænsefeltet mellem sundhedsvidenskab og humaniora/samfundsvidenskab. Tidsskriftets målsætning er at fungere som et forum, hvor disse fag kan mødes og inspirere hinanden – epistemologisk, metodisk og teoretisk – i forskellige forskningssammenhænge. Tidsskriftet formidler den debat og teoretiske udvikling, der foregår i de voksende faglige samarbejds- og forskningsinitiativer, der udspringer af dette grænsefelt. Tidsskriftet henvender sig til alle med interesse for forskning i sygdom og samfund og i særlig grad til sundhedsmedarbejdere i forsknings- og undervisningssammenhæng med forbindelse til tværfaglige miljøer.*

## Aims and scopes

*The Journal for Research in Sickness and Society is an interdisciplinary journal which has a theoretical background in medical anthropology. The aim and purpose of the journal is to promote and develop research in the borderland between the health sciences and the humanities/the social sciences. The goal of the journal is to function as a forum in which these disciplines may meet and inspire each other – epistemologically, methodologically and theoretically. The journal conveys the debate and theoretical development which takes place in the growing collaboration and research initiatives emerging from this borderland. The journal addresses all with an interest in research in sickness and society and especially health professionals working with education and/or research in interdisciplinary institutions.*

# Indhold

*Claus Bossen, Maja Hojer Bruun & Sune Hannibal Holm*

Kunstig intelligens og sundhed: En introduktion 5-18

*Maria Bach Nielsen, Tom Børsen & Casper Knudsen*

Before Deployment: Anticipatory Infrastructuring and Early AI Integration in Acute Stroke-ready MRI Workflow 19-47

*Michal Frumer, Christina Egelund Antonsen & Maja Hojer Bruun*

AI som beslutningsstøtte til fraktur i skadestuen? Billeddiagnostik, tidslighed og patientforløb i afprøvningen af fraktur-AI 48-75

*Iben Mundbjerg Gjødsbøl & Mette Nordahl Svendsen*

Personlige prognoser: Forudsigelser af fremtiden med kunstig intelligens (AI) i klinisk praksis 76-99

*Mari Serine Kannelønning & Miria Grisot*

Collective efforts for the introduction of AI in healthcare: An empirical study of two Norwegian initiatives 100-123

*Victor Vadmand Jensen & Jens Christian Bjerring*

Between Theoretical and Empirical Ethics of Healthcare Artificial Intelligence: The Case of Autonomy in Breast Cancer Screening 124-147

*Camilla Brændstrup Laursen*

Homes, Homeliness, and Otherness in Medical Anthropology “at Home” 148-169

Forfatterliste 169-173

Retningslinjer for forfattere 174-178

## *Forsidetekst*

Maleri: *Mødes de?*

Kunstneren:

”Mødes de?” er titlen på maleriet af de to personer. Man får indtryk af, at der er et møde, men er de på vej væk fra hinanden, eller er de på vej til at være sammen. Den røde farve på den ene person er varm, mens den anden person er en kold blå, så vil det være muligt?

Billedet er en del af en udstilling i Galleri Gallo, Aarhus, som hedder Nærvær-Samvær-Fravær. Den røde tråd i den udstilling har været at undersøge vores relationer og den måde, vi agerer med hinanden.

Jeg har et værksted i Aarhus C, Eugen Warmingsvej 35 kld., hvor man er velkommen til et besøg efter aftale.

Se også Instagram: [elisabethbuhrkal](#)

*Elisabeth Buhrkal*

# Kunstig intelligens og sundhed

## *En introduktion*

Claus Bossen<sup>1</sup>

Maja Hojer Bruun<sup>2</sup>

Sune Hannibal Holm<sup>3</sup>

---

<sup>1</sup>Informationsvidenskab, Aarhus Universitet

<sup>2</sup>Pædagogisk Antropologi, Aarhus Universitet

<sup>3</sup>Sektion for Forbrug, Bioetik og Regulering, Københavns Universitet

Bossen, Claus; Bruun, Maja Hojer & Holm, Sune Hannibal. 2026. Kunstig intelligens og sundhed: En introduktion. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 5-18. DOI: 10.7146/tfss.v25i44.168732

Det er næppe undgået nogens opmærksomhed, at kunstig intelligens – eller Artificial Intelligence (AI) – er kommet til det danske sundhedsvæsen. Kunstig intelligens er siden 2010'erne gået ind i sin 'tredje sommer', og det er sidenhen kun blevet varmere. Den første varme interesse – heraf metaforen 'sommer' – kom i årene lige efter Anden Verdenskrig, men efter en ti års tid kølnedes interessen, da feltet syntes at stå i stampe og gik ind i en 'vinter'-periode. Større computerkraft førte til den anden sommer fra midt-1970'erne til midt-1980'erne igen efterfulgt af en vinter, hvor interessen for kunstig intelligens var minimal (Kautz, 2022). Den tredje sommer startede først i forskningskredse og store IT-firmaer (Meta, Google, m.fl.) efterfulgt af store private investeringer og offentlige forskningsbevillinger. De store sprogmodeller (Large Language Models) udvikledes i 2010'erne med BERT, ChatGPT 1 og 2, BART og Llama m.fl., og er blevet mere og mere avancerede. Det store gennembrud for kunstig intelligens i offentligheden kom med lanceringen i 2022 af frit tilgængelige platforme for såkaldte generative systemer

såsom ChatGPT, der producerer tekst, og Dall-E, der producerer billeder. I løbet af ganske få år er kunstig intelligens kommet ind i mange forskellige aspekter af vores arbejds- og fritidsliv. Visioner om, hvordan kunstig intelligens kan forbedre stort set alle områder af vores samfund, lanceres vedvarende med stor entusiasme.

## Kunstig intelligens i sundhedssektoren

I sundhedssektoren er der en række rationaler og begrundelser for at introducere AI, som går igen i den offentlige debat på overordnet politisk plan og i konkrete praktisk-organisatoriske beslutninger. De fortjener at blive underkastet nøjere eftersyn: Et hyppigt fremført argument knytter sig til det stigende pres på sundhedsvæsenet, som blandt andet skyldes en aldrende befolkning, en stigning i kroniske sygdomme samt mangel på både sundhedspersonale og speciallæger (Danske Regioner, 2024; Regeringen, KL og Danske Regioner, 2025). I denne optik fremstilles AI som en mulig løsning, der kan afhjælpe systemets belastning ved at overtage og automatisere mere "rutineprægede" opgaver samt bidrage til en mere effektiv prioritering af patienter. Oplagte spørgsmål for en samfundsforsker i denne sammenhæng er at problematisere, hvilke værdier, som knytter sig til prioriteringerne, hvad der forstås ved "rutinearbejde" eller "tidsbesparelse", og hvorvidt der i det arbejde, der søges effektiviseret, også indgår værdifulde men usynlige eller oversete arbejdsopgaver.

Et andet gennemgående rationale handler om bedre ressourceudnyttelse. Her fremstår kunstig intelligens ofte som et løfte om et mere smidigt og produktivt sundhedsvæsen, hvor ressourcerne anvendes mere optimalt. Ligesom Klaus Hoyer (2023) har efterspurgt empirisk belæg for faktisk opnåede fordele gennem dataficering, kan vi i dag spørge, i hvilket omfang de løfter, som AI kommer med, bliver indfriet og dokumenteret i praksis. Derudover fremhæves potentialet for forbedring af diagnostik og behandling som et væsentligt argument. AI-systemer beskrives her som teknologier, der kan identificere komplekse mønstre i store datamængder og dermed muligvis opdage diagnoser eller sammenhænge, som kan være vanskelige for det menneskelige blik at få øje på. I forlængelse heraf knyttes også forventninger om mere individualiseret behandling og udvikling af såkaldt personlig medicin, hvor behandlingen i højere grad tilpasses den enkelte patients biologiske og genetiske profil. Et yderligere aspekt handler om udnyttelsen af de store mængder sundhedsdata, som er blevet tilgængelige gennem tidligere års digitalisering og dataficering af sundhedssystemet. Disse data fremstår som en ressource, der endnu ikke er fuldt udfoldet, og som AI-teknologier potentielt kan bringe i spil på nye måder (Tupasela et al., 2020).

Politiske og strategiske hensyn spiller også en væsentlig rolle, når der i dag afprøves og investeres i kunstig intelligens. I Danmark har kunstig intelligens fået statslig bevågenhed med "Strategi for digital sundhed 2018-2022" (siden forlænget til 2024) (Sundhedsdatastyrelsen, 2018), der primært fokuserede på digitalisering bredt, men også pegede på kunstig intelligens som drivkraften i 'datadrevne arbejdsgange' og 'tidlig opsporing af sygdom og beslutningsstøtte'. Siden har regeringen for eksempel lanceret "National strategi for kunstig intelligens" (Digitaliseringsstyrelsen, 2019) som blandt andet finansierede og igangsatte en række 'signaturprojekter', der skulle undersøge potentialerne for kunstig intelligens i den offentlige sektor, herunder sundhedsvæsenet. Interessen for kunstig intelligens er bred inden for sundhedsvæsenet og hos private firmaer, og der er i disse år en lang række projekter under udvikling.

I nogle sammenhænge kan ambitionen om at investere i AI fremstå som et mål i sig selv, hvor Danmark – eller en konkret region, et konkret hospital inden for regionen eller et sundhedsfagligt speciale – skal være "forrest i udviklingen". Endelig kan vi notere, at nogle af engagementerne i medicinsk og teknologisk forskning er drevet af en grundlæggende nysgerrighed i forhold til nye teknologiers muligheder og grænser. For øjeblikket udvikles og afprøves en lang række nye metoder og redskaber, som måske aldrig ender med at lykkes eller udbredes. Sammenfattende er den faktiske ibrugtagning af forskellige systemer med kunstig intelligens i klinikken begrænset trods den store begejstring og offentlige støtte, ikke kun i Danmark, men også internationalt (Zajac et al., 2023). Udfordringerne med valide data, udvikling af modeller, integrering i arbejdsgange, etik samt juridiske forhold angående blandt andet GDPR og beskyttelse af personoplysninger har vist sig komplicerede (Zajac et al., 2023; Deloitte, 2024), og der er kritiske røster, der påpeger risici ved ændrede roller for patienter og klinikere samt ved diagnosticering og behandling (Svensson & Jotterand, 2022). I denne sammenhæng er det vigtigt at understrege, ligesom Suchman (2023) gør det, at kunstig intelligens ikke er én entydig og historieløs ting, men en lang række forskellige tekniske "løsningsforslag" på problemer, som i sig selv kunne artikuleres på helt andre måder. Det er derfor vigtigt, at vi fastholder et fokus på AI-teknologiers lokale kontekster og forskelligartede politiske, historiske, sociale og kulturelle effekter.

Udviklingen af AI-systemer til beslutningsstøtte har givet anledning til en lang række etiske overvejelser. En af de mest centrale debatter vedrører risici for at AI-systemer ikke fungerer lige godt for forskellige sociale grupper. Det vel nok mest omdiskuterede eksempel på denne såkaldte bias og fairness problematik

er COMPAS-algoritmen, der anvendes i retsvæsenet i en række stater i USA. En analyse af COMPAS-algoritmen indikerede, at den ikke var lige god til at forudsige risiko for recidivisme for grupper af forskellige racer og køn (Angwin et al., 2016).

Flere studier har dokumenteret, hvordan de datasæt, der anvendes til at træne og teste algoritmer, ofte har en underrepræsentation af relevante grupper, og derfor fungerer bedst på den eller de grupper, der udgør majoriteten i de anvendte datasæt. Wen et al. (2024) påviser for eksempel stor ubalance i datasæt med hudlæsningsbilleder med hensyn til etnicitet, hudtoner og geografisk oprindelse, hvilket kan føre til stor ulighed i algoritmens nøjagtighed for eksempelvis forskellige etniske grupper. Perez (2019) giver adskillige eksempler på, hvordan kvinder historisk har været underrepræsenteret i datasæt og i designet af mange former for samfundsinfrastruktur, lige fra trafikplanlægning til algoritmiske systemer. Efterhånden som algoritmer bliver udbredte og bruges til at træffe beslutninger om allokering af ressourcer og muligheder til millioner af mennesker, bliver det kun endnu mere relevant at sikre klarhed om, hvilke datasæt sundhedsalgoritmer trænes og testes på.

Problematiske algoritmisk bias kan også opstå i forbindelse med anvendelse af såkaldte "proxier". Obermeyer et al. (2019) påviser eksempelvis, hvordan en kommerciel algoritme, som bliver bredt anvendt i det amerikanske sundhæsvæsen til at afgøre, hvilke patienter, der vil få mest gavn af særlige sundhedsprogrammer, vurderer sorte amerikanere til at være mindre alvorligt syge end hvide amerikanere, selvom de faktisk er lige syge. Den pågældende bias opstår, fordi algoritmen forudsiger sundhedsomkostninger som proxy for sygdom, men da den sociale stratifikation i høj grad overlapper med meget ulige adgang til sundhed og pleje for sorte og hvide, betyder det, at der bruges færre penge på sorte patienter end hvide patienter, når de faktisk er lige alvorligt syge.

En anden debat, der har ledt til en større etisk diskussion, vedrører behovet for at kunne forklare, hvorfor en AI kommer med en bestemt forudsigelse vedrørende et individ. For så vidt at sundhedspersonale forventes at anvende AI-forudsigelser om patienter som beslutningsstøtte, forekommer det vigtigt fra såvel personale- som patientperspektiv, at sundhedspersonale kan forklare, hvorfor AI'en kom med den pågældende forudsigelse om patienten. Dels er det relevant for personalet at kunne forklare en forudsigelse, der ligger til grund for en beslutning, til den berørte patient, og dels er det vigtigt for en ansvarlig beslutningstager at kunne validere en forudsigelse ved at have indblik i, hvordan AI-systemet kom frem til den.

I den forbindelse er det en påtrængende udfordring, at mange AI-systemer er så komplekse, at det ikke er muligt for mennesker at gennemskue, hvordan systemet

når frem til en forudsigtelse fra et input. Dette omtales almindeligvis som black box-effekten (Latour, 1999, s. 304), hvor det kun er input og output, der er synligt, mens maskinens interne processer er lukket af som i en sort boks. I de seneste år har udviklere og producenter af AI-systemer forsøgt at tage den kritik til sig, at fagprofessionelle har brug for at vide, hvordan anbefalinger, risikoscoringer og prioriteringer er blevet til, og har udviklet systemer med såkaldt forklarlighed. Et voksende forskningsfelt vedrørende Explainable AI eller XAI (på dansk: forklarlighed) arbejder med at udvikle metoder, der kan give en form for forklaring på AI-forudsigtelser (Baredo Arrietta et al., 2020). Det er imidlertid omdiskuteret, om disse metoder er troværdige (Ghassemi et al., 2021; Holm 2023). Forklarlighed kræver ofte tilpasning til specifikke organisationer og konkrete formål, hvilket kræver gode og grundige implementeringsprocesser. Igen er det ikke sikkert, at alle redskaber med kunstig intelligens ender med at blive anvendelige i praksis, så der er brug for en stor grad af kritisk tænkning, faglig ekspertise og åbenhed for at kunne fravælge fejlslagne eller uegnede teknologier.

Med henblik på at fremme etisk forsvarlig udvikling og anvendelse af AI-systemer har Europa-Kommissionen udgivet sine retningslinjer for Trustworthy AI (HLEG, 2019). Retningslinjerne for Trustworthy AI identificerer fire etiske principper, der bør vejlede udviklingen og brugen af AI-systemer: (i) Respekt for menneskelig autonomi, (ii) Forebyggelse af skade, (iii) Fairness og (iv) Forklarlighed. For eksempel fastslår princippet om respekt for menneskelig autonomi, som bliver diskuteret i Jensen og Bjerrings artikel i dette temanummer, at "Mennesker, der interagerer med AI-systemer, skal være i stand til at bevare fuld og effektiv selvbestemmelse" (HLEG, 2019). I et forsøg på at sikre, at de generelle principper også finder anvendelse i praksis, er der udarbejdet en liste, der kan hjælpe udviklere og brugere med at vurdere konkrete værktøjer i forhold til de generelle principper (ALTAI, 2020). Listen er et forsøg på at gøre det nemmere for udviklere at forstå, hvad Trustworthy AI er, og hjælpe dem til at identificere hvilke risici et AI-system kan generere, og hvordan man minimerer disse risici under hensyntagen til de forventede fordele ved AI-systemet. Det er uklart, i hvilken udstrækning ALTAI bliver anvendt, men enkelte forskningsgrupper har forsøgt at gennemføre systematiske analyser af konkrete anvendelser af AI-systemer inden for sundhedsvæsenet i lyset af EU's retningslinjer (Amann et al., 2022; Vetter et al., 2023). Samlet set forekommer der at være behov for endnu større fokus på, hvordan generelle etiske retningslinjer skal fortolkes og anvendes i praksis.

I foråret 2024 vedtog Europa-Parlamentet og EU-rådet "Forordningen vedrørende kunstig intelligens", som klassificerer AI-anvendelser efter deres risiko for at forvolde skade i fire niveauer: uacceptabelt, højt, begrænset og minimalt (Den

Europæiske Union, 2024). Forordningens krav til anvendelse af kunstig intelligens skal indføres gradvist frem til 2028. Erhvervslivets lobbyorganisationer presser dog for øjeblikket på for at indføre længere overgangsordninger og lempeligere krav, ofte med henvisning til EU's konkurrencedygtighed i det globale kapløb om at udvikle AI-teknologier. Sådanne forenklinger og lempelser af AI-forordningen er indeholdt i Europa-Kommissionens forslag fra november 2025 om en Digital Omnibus, også kaldet AI Omnibus (Europa-Kommissionen 2025). Udfaldet af disse verserende politiske diskussioner og beslutninger vil få stor betydning for kunstig intelligens i sundhedssektoren i fremtiden.

## Artiklerne i dette temanummer

Dette temanummers fem artikler analyserer forskellige former for AI og viser, hvordan lokale kontekster og processer er med til at forme, hvad AI bliver til i praksis, jvf. Suchman (2023) ovenfor. Alle artiklerne er funderet i empiriske studier, hvoraf de tre artikler fokuserer på implementering og testning af AI i kliniske kontekster: Nielsen, Børsen og Knudsen har rettet blikket mod den tidlige implementering af det AI-baserede beslutningsstøttesystem Apollo, som skal hjælpe neurologer og radiologer med at sortere og prioritere patienter med slagtilfælde ved at analysere på tværs af forskellige MR-scanninger af hjernen. Frumer, Antonsen og Bruun har undersøgt testningen af fraktur-AI, som markerer knoglebrud og ledeskred på røntgenbilleder på to danske skadestuer. Endeligt analyserer Gjødsbøl og Svendsen CARDIA<sup>HD</sup>-algoritmen, et risikovurderingsværktøj trænet på omfattende sundhedsdata, herunder diagnose- og procedurekoder, laboratorieresultater og kliniske målinger, til at forudsige overlevelsesprognosen for patienter med akut iskæmisk hjertesygdom, som er implementeret på forsøgsbasis blandt kardiologer. Den fjerde artikel af Kannelønning og Grisot har i stedet fokus på, hvorledes offentlige diskussioner af AI i Norge udfolder sig i to fora, der havde til formål kollektivt at udvikle viden om problematikker og løsninger omkring sundheds-AI, mens den femte artikel af Jensen og Bjerring med udgangspunkt i en empirisk case tager fat på en filosofisk diskussion: Kan et etisk krav om, at autonomi i beslutninger omkring patienter alene ligger hos klinikerne opretholdes i praksis, når AI indgår i processen?

Nielsen, Børsen og Knudsen beskriver og analyserer i deres artikel *“Before Deployment: Anticipatory Infrastructuring and Early AI Integration in Acute Stroke-ready MRI Workflow”* det arbejde, der skal gøres, før AI kan implementeres i en klinisk hverdag. De kalder dette ‘anticipatorisk infrastrukturering’ ud fra en tilgang,

der tilsiger, at infrastrukturer, når de skal opbygges, kræver et arbejde med at sætte forskellige sociale og tekniske elementer sammen. Selv når de er etablerede, kræver de vedligeholdelse, reparationer og udbygning, så infrastrukturen er konstant i proces. Det anticipatoriske element peger så specifikt på det arbejde, der er nødvendigt for at kunne få det tekniske og det sociale til at passe sammen. I deres case omkring MR-scanning af akut indlagte patienter med blodprop i hjernen handler det blandt andet om, hvad AI-systemets rolle skal være: Skal den rådgive eller må den tage beslutninger, og selv hvis AI'en kun rådgiver, er der så et øget ansvar for lægerne? Hvad sker der, hvis de ikke er enige med AI'en, og det senere viser sig, at der faktisk var en komplikation? Hvordan skal AI-systemet integreres i en arbejdsproces, der i forvejen er hurtig, uden at flowet afbrydes? Skal en rød knap lyse, hvis AI'en finder noget kritisk? Artiklen peger dels på, at AI ikke gnidningsløst glider ind i de eksisterende IT-systemer og arbejdsprocesser, og dels på de specifikke problemstillinger der rejser sig, når læger, teknikere, projektledere og -deltagere skal forberede brugen af AI i klinikken.

Med udgangspunkt i etnografisk feltarbejde og interviews på to danske skadestuer, analyserer Frumer, Antonsen og Bruun i et tredje empirisk casestudie, *“AI som beslutningsstøtte til fraktur i skadestuen? Billeddiagnostik, tidslighed og patientforløb i afprøvningen af fraktur-AI”*, hvordan fraktur-AI indgår i klinisk praksis og formes af skadestuernes organisatoriske rytmer, samarbejder mellem klinikere og politiske forventninger. Artiklen argumenterer for, at fraktur-AI ikke blot fungerer som et teknisk hjælpemiddel, men må forstås som del af et socioteknisk system, hvor teknologiens betydning afhænger af den konkrete kliniske kontekst. Analysen viser for det første, hvordan fraktur-AI bidrager til en “punktuering af tid”, hvor komplekse diagnostiske processer reduceres til binære vurderinger af brud eller ikke-brud. Sammenholdt med øvrige kliniske praksisser kan dette skabe tryghed og afsæt for kliniske beslutninger hos yngre læger, men opleves samtidig som forstyrrende og kontekstløst af mere erfarne klinikere. For det andet undersøges forholdet mellem AI, tempo og ventetid. Fraktur-AI legitimeres politisk gennem forestillinger om effektivisering og hurtigere patientflow, men empirien viser, at diagnostik og behandling forbliver afhængige af klinisk erfaring, konferering og tværfagligt samarbejde. Endelig viser artiklen, hvordan afprøvningen bæres af et stærkt fremtidsløfte om, at AI med tiden vil blive mere præcis. Selvom mange læger oplever teknologien som mangelfuld i nutiden, opretholdes legitimiteten gennem forventninger om fremtidige forbedringer. Artiklen bidrager dermed med et kritisk perspektiv på AI som en teknologi, der gør visse fremtidsforventninger styrende for nutidens kliniske praksisser og organisatoriske beslutninger.

Gjødtsbøl og Svendsen undersøger i deres artikel *“Personlige prognoser: Forudsigelser af fremtiden med kunstig intelligens (AI) i klinisk praksis”*, hvordan kardiologer i klinisk praksis forstår AI-beregnete overlevelsesprognoser for patienter med akut iskæmisk hjertesygdom. Disse prognoser leveres af den såkaldte CARDIA<sup>IHD</sup>-algoritme i form af indplacering af individuelle patienter i en af fem risikoklasser i forhold til risikoen for at dø inden for et år. Risikovurderingsværktøjet blev implementeret på forsøgsbasis i det elektroniske patientjournalssystem Sundhedsplatformen i 2023 ved hospitaler i Region Hovedstaden og Region Sjælland, hvor kardiologerne er ansat. Gjødtsbøl og Svendsen viser i deres artikel, hvordan kardiologer arbejder med at vurdere og indarbejde CARDIA<sup>IHD</sup>-algoritmens anbefalinger i deres kliniske praksis: For det første mangler algoritmen ofte væsentlige data, enten fordi data ikke opdateres i realtid eller fordi patientens sygehistorie ikke ligger på Sundhedsplatformen. Som på så mange andre områder må lægerne selv fortolke data og tilføje social, kulturel og historisk kontekst til algoritmens klassificering af patienterne. For det andet er kardiologer vant til at arbejde med risikovurderingsredskaber, der bygger på kendte årsagssammenhænge. Derfor kan AI-algoritmers risikovurderinger, der blot sker på baggrund af korrelationer i data, virke forvirrende og føre til, at algoritmens forudsigelser ikke anvendes i praksis. Sidst men ikke mindst opererer forskellige dele af sundhedsvæsenet med forskellige tidshorisonter for patienters indlæggelse og behandling. Højt specialiserede kardiologer vil ofte kun se patienterne op til en eventuel operation og efter, at den er overstået, hvorefter patienterne overføres til andre afdelinger. Deres tidshorisonter er således som regel kortere end sandsynligheden for overlevelse inden for et år, som algoritmen beregner. Prognoser som dem, risikovurderingsværktøjer leverer, skal derfor altid ses i relation til den tidlige og rumlige organisering af de patientforløb, som de indgår i.

Kannelønning og Grisot peger i deres artikel *“Collective efforts for the introduction of AI in healthcare: An empirical study of two Norwegian initiatives”* på en anden form for anticipatorisk infrastrukturering, hvor det ikke som med Nielsen m.fl. handler om en konkret case, men mere generelt om, hvordan vi som samfund kollektivt kan opbygge viden om kunstig intelligens. Kannelønning og Grisot sætter fokus på, hvordan to forskellige fora i Norge har arbejdet med at skabe viden om AI i forbindelse med indførelsen af AI i sundhedsvæsenet. De tager fat på to initiativer: Dels et forum etableret af regeringen for at udarbejde en vision for AI, og dels et uformelt netværk af folk, der arbejder professionelt med AI indenfor sundhed. Deres udgangspunkt er, at indførelsen af AI i sundhedsvæsenet skal være et kollektivt anliggende, idet det må være en kollektiv proces at finde frem

til bæredygtige og etisk gode AI-løsninger og -integrationer. Processerne med at skabe viden i sådanne initiativer kan, skriver de, orkestreres via konsensus, dominans ovenfra, eller via en hybrid af disse. De viser således, hvordan der i begge tilfælde var hybride elementer på spil, og peger på vigtigheden af at kunne håndtere forskellige interesser og målsætninger samt på vigtigheden af at være bevidst om forskellige orkestrerings-modeller. Kollektive processer er, argumenterer de, veletablerede i de skandinaviske lande og vigtige i forbindelse med udvikling, implementering og skalering af AI, således at AI-teknologierne bliver genstand for offentlig debat.

I de tre studier af implementering af AI-systemer af henholdsvis Nielsen, Børsen og Knudsen, Frumer, Antonsen og Bruun og Gjødsbøl og Svendsen, understreges det, at AI-redskaberne kan udgøre en form for beslutningsstøtte, men selve beslutningen og ansvaret for den ligger hos lægerne. Netop dette spørgsmål om ansvar og etik diskuteres i artiklen *“Imellem Teoretisk og Empirisk Etik for Kunstig Intelligens i Sundhedsvæsenet: En Case om Autonomi i Brystkræftscreening”* af Jensen og Bjerring. Med udgangspunkt i filosofisk etik bidrager forfatterne med en tiltrængt overvejelse over forholdet mellem generelle etiske og teoretisk begrundede principper og såkaldt empirisk etik, der fokuserer på etiske normer som kontekstbestemte i den kliniske praksis. I deres artikel diskuterer forfatterne forholdet mellem generelle etiske retningslinjer og principper såsom principperne præsenteret i EU’s retningslinjer og den konkrete anvendelse af AI-systemer i sundhedsvæsenet. Som illustration af den generelle pointe viser Jensen og Bjerring, hvordan teoretiske overvejelser, der klassificerer en anvendelse af AI som uetisk, bliver modsagt af de etiske normer, der formes af den kliniske kontekst. Som illustration af deres pointe analyserer Jensen og Bjerring, hvorledes princippet om respekt for menneskelig autonomi forhandles, når der implementeres AI-beslutningsstøttesystemer i relation til AI i brystkræftscreening. Fokus for Jensen og Bjerrings diskussion er det bredt accepterede syn på AI-autonomi, *“The Non-Autonomous AI Principle”*, ifølge hvilket AI, der anvendes til beslutningsstøtte i sundhedsvæsenet, aldrig bør have samme autoritet som klinikere og altid bør kunne afvises af menneskelige beslutningstagere. Gennem en analyse af anvendelsen af AI i brystkræftscreening argumenterer de for, at autonom AI dog kan accepteres i klinisk praksis på trods af, at det afvises principielt. På den baggrund diskuterer Jensen og Bjerring behovet og mulighederne for at skabe etiske retningslinjer for anvendelse af AI i klinisk praksis, der ikke tager udgangspunkt i generelle principper, men derimod i etik som noget, der etableres i en konkret klinisk kontekst.

## Tværgående temaer og pointer

Artiklerne rummer en række tværgående temaer og pointer, som vi kort vil komme ind på her. Én tematik er AI som støj. Både Frumer, Antonsen og Bruun og Gjødsbøl og Svendsen peger på, at læger kan opfatte visse oplysninger fra AI-algoritmer som 'støj' i forhold til deres vante arbejdspraksis, i dette tilfælde de dele af algoritmens såkaldte 'forklarende faktorer', som ikke umiddelbart giver mening, og som de ikke kan handle på. Spørgsmålet er, om AI'ens resultater, vurderinger eller anvisninger opfattes som en hjælp til det kliniske arbejde, som en form for støj, der filtreres fra, så AI'en faktisk ikke anvendes, eller om den skaber nye eller yderligere former for opgaver, ansvar og forpligtelser, som udgør et ekstra arbejde for klinikerne. I Nielsen, Børsen og Knudsens artikel er det et vigtigt opmærksomhedspunkt i den tidligere implementering af Apollo, at systemet ikke må føre til støj for klinikerne, men at Apollos indflydelse begrænses til relevante oplysninger, at systemets tempo og timing tilpasses eksisterende arbejdsgange, og at systemet ikke tiltrækker sig mere af klinikernes opmærksomhed end højest nødvendigt. Apollos markeringer kan let udvikle sig fra støtte til støj, hvor visuelle alarmer og tidsmæssige forskydninger forstyrrer lægernes rytmer for opmærksomhed, prioritering og ansvar i det billeddiagnostisk arbejde.

Hvordan AI-systemer (risikerer at) strukturere(r) tid og tidslighed på nye måder er et andet tværgående tema. Gjødsbøl og Svendsen beskriver, at CARDIA<sup>IHD</sup>-algoritmen introducerer en ny kategori af personlige prognoser med en ny tidshorisont på et år, som ikke er relevant for alle hospitalsafdelinger, mens fraktur-algoritmer, som beksrevet af Frumer, Antonsen og Bruun, samler fokus på ét punkt i tid, fastslåelsen af brud eller ej, hvor unge læger har brug for meget mere information og ikke mindst sparring med erfarne kollegaer for at behandle forskellige typer brud. Den vellykkede integration eller infrastrukturering af Apollo afhænger af, at systemets tempo passer til det kliniske tempo; ellers skabes forsinkelser, usikkerhed og friktion. En vellykket implementering af forskellige former for kunstig intelligens i sundhedsvæsenet afhænger ikke alene af en vurdering af teknologierne i sig selv, deres præcision eller effektivitet, men også deres samspil med organisationernes eksisterende arbejdsgange, patientforløb, kommunikationsformer og informationsstrømme. Organisationsetnografen Barley (1986) brugte begrebet 'strukturering' til at vise, hvordan teknologi omformer sociale og organisatoriske relationer. Introduktionen af nye teknologier i en organisation kan derfor blive en god anledning til at vurdere og bevidst forbedre organisationens strukturering af arbejdsprocesserne.

Hvis vi løfter blikket lidt op fra de fem artikler, rejser der sig andre spørgsmål. Et spørgsmål er for eksempel, om AI reelt giver værdi. Det er en kendt pointe

fra indførelsen af nye sundhedsinformationssystemer såsom den elektroniske patientjournal (EPJ) (Vikkelsø, 2005), at de ikke nødvendigvis gør organisationerne *mere effektive* eller *fjerner* overflødige arbejdsopgaver, men at de snarere bidrager til, at arbejdsopgaver, opmærksomhed og risiko *omfordeles*, så nye opgaver, ansvarsområder og risici opstår, hvor andre forsvinder. Her udgør AI ikke nogen undtagelse, selvom nogle studier viser, at AI kan føre til tidsbesparelse for klinikere for eksempel indenfor mammografi-screening (Lauritsen et al., 2022), og selvom der i skrivende stund (maj 2026) er store forventninger til, at AI-baserede samtale-til-tekst-løsninger kan spare mange læge- og lægesekretær timer. Indtil videre virker det til, at AI-systemer kan komme til at virke i klinisk praksis, hvis de er udviklet til konkrete, målrettede opgaver, eksempelvis automatisering af begrænsede, velkendte arbejdsprocesser og samkøringer af dataformer, som ligger i forlængelse af eksisterende journalssystemer og informationsflows i EPJ frem for at udgøre et alternativ til eller erstatte dem.

Samtidig er der imidlertid endnu ikke så mange AI-redskaber, som er klinisk validerede (se dog Lauritsen et al., 2022), og der er blandt praktikere en vis skepsis over for AI-hype forstået på den måde, at AI tillægges en særlig værdi og gives særstatus, blot fordi teknologierne har nyhedens interesse. Fra samfundsvidenskabelige forskere, som arbejder tæt sammen med klinikere, der har erfaring med implementering af kunstig intelligens i radiologi og andre områder, lyder det, at AI-implementeringer bør involvere de sundhedsprofessionelle, der skal bruge teknologierne, og finde sted uden eksternt politisk og kommercielt pres (Kristensen, 2025), og at nytteværdien af AI må række ud over spørgsmål om tidsbesparelser og for eksempel også kan forbedre arbejdskvaliteten og arbejdsmiljøet i sundhedsvæsenet (Jensen, 2025). I forhold til konkrete AI-projekter kunne samfundsforskere spørge, om det er etisk og økonomisk forsvarligt at anvende omfattende ressourcer til forskning og udvikling af nye teknologier, hvis værdi i sidste ende afhænger af vellykkede implementeringsprocesser, idet disse ressourcer kunne være anvendt bedre på anden vis. Der findes cost-benefit analyser, men de er desværre sjældne (se dog Rosen et al., 2025, Cresswell et al., 2025). Ydermere er der vigtige spørgsmål om, hvilken rolle eksempelvis diagnostiske algoritmer bør spille i klinisk beslutningstagning, og hvorledes de kan implementeres i beslutningsprocesser på en måde, der er gavnlig for samfundet og for patienter. Disse problemstillinger behandles i dette temanummer af Nielsen, Børsen og Knudsen samt af Jensen og Bjerring.

Kunstig intelligens stiller os overfor en række spørgsmål i krydsfeltet mellem medicin, datalogi og forskellige samfundsvidenskabelige og humanistiske tilgange til sundhed: Hvad sker der med samspillet med borgere og sundhedsvæsenet

ved indføringen af kunstig intelligens? Hvordan kan kunstig intelligens' anbefalinger gøres transparente, retfærdige, forklarlige og forståelige? Hvordan sikrer vi, at dens indførelse kommer patienter og sundhedsprofessionelle til gode og ikke bliver (endnu) en fremmedgørende effektiviseringsteknologi? Her kan samfundsforskere bidrage med analyser af, hvordan teknologier omformer professionelle relationer, koordinationsformer og organisatoriske praksisser i konkrete kliniske sammenhænge.

Med dette temanummer af Tidsskrift for Forskning i Sygdom og Samfund ønsker vi at bidrage til en debat og øget viden om hvad kunstig intelligens er og kan, og hvad den rejser af forskellige problemstillinger. Særnummeret indeholder fem artikler, der på hver deres måde præsenterer indsigtfulde analyser af kunstig intelligens og sundhed.

God fornøjelse med temanummeret,

Claus Bossen, Maja Hojer Bruun, Sune Hannibal Holm

## Referencer

- Amann, J., Vetter, D., Blomberg, S. N., Christensen, H. C., Coffee, M., Gerke, S., et al. (2022). To explain or not to explain? –Artificial intelligence explainability in clinical decision support systems. *PLOS Digital Health*, 1(2), e0000016. <https://doi.org/10.1371/journal.pdig.0000016>
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). *Machine bias: There's software used across the country to predict future criminals and it's biased against blacks*. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Barley, S. R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly*, 31(1), 78–108. <https://doi.org/10.2307/2392767>
- Cresswell, K., Williams, R., Dungey, S., Anderson, S., Bernabeu, M. O., Mozaffar, H., Yang, X., Sai, V., Bea, S., & Eason, S. (2025). A mixed methods formative evaluation of the United Kingdom National Health Service Artificial Intelligence Lab. *Digital Medicine*, 8(1), 448. <https://doi.org/10.1038/s41746-025-01835-2>
- Danske Regioner. (2024). *Udspil: Om kunstig intelligens. Vi skal skabe et godt og effektivt sundhedsvæsen ved brug af kunstig intelligens*.
- Deloitte. (2024). *Evaluering af AI-signaturprojekter*. Digitaliseringsstyrelsen, Kommunernes Landsforening og Danske Regioner.

- Den Europæiske Union. (2024). *Europa-Parlamentets og Rådets Forordning om harmoniserede regler for kunstig intelligens og om ændring af forordning (EF) nr. 300/2008, (EU) nr. 167/2013, (EU) nr. 168/2013, (EU) 2018/858, (EU) 2018/1139 og (EU) 2019/2144 samt direktiv 2014/90/EU, (EU) 2016/797 og (EU) 2020/1828 (forordningen om kunstig intelligens)*. <https://eur-lex.europa.eu/legal-content/DA/TXT/?uri=CELEX:32024R1689>
- Digitaliseringsstyrelsen. (2019). *National strategi for kunstig intelligens*. [https://digst.dk/media/19302/national\\_strategi\\_for\\_kunstig\\_intelligens\\_final.pdf](https://digst.dk/media/19302/national_strategi_for_kunstig_intelligens_final.pdf)
- Europa-Kommissionen. (2025). *Forslag til Europa-Parlamentets og Rådets Forordning om ændring af forordning (EU) 2024/1689 og (EU) 2018/1139 for så vidt angår forenkling af gennemførelsen af harmoniserede regler om kunstig intelligens (digital omnibus om kunstig intelligens)*.
- Ghassemi, M., Oakden-Rayner, L., & Beam, A. (2021). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11), e745–e750. [https://doi.org/10.1016/S2589-7500\(21\)00208-9](https://doi.org/10.1016/S2589-7500(21)00208-9)
- High-Level Expert Group on Artificial Intelligence. (2019). *Ethics guidelines for trustworthy AI*. European Commission. <https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html>
- Hoeyer, K. (2023). *Data paradoxes: The politics of intensified data sourcing in contemporary healthcare*. MIT Press.
- Holm, S. (2023). On the justified use of AI decision support in evidence-based medicine: Validity, explainability, and responsibility. *Cambridge Quarterly of Healthcare Ethics*, 33(3), 1–7. <https://doi.org/10.1017/S0963180123000324>
- Jensen, V. V. (2025). Forsker: Skal AI transformere sundhedsvæsenet, må vi se ud over tidsbesparelser. *Altinget*. <https://www.alinget.dk/digital/artikel/forsker-skal-ai-transformere-sundhedsvaesenet-maa-vi-se-ud-over-tidsbesparelser>
- Kautz, H. (2022). The third AI summer. *AI Magazine*, 43(1), 105–125. <https://doi.org/10.1002/aaai.12036>
- Kristensen, D. B. (2025, January 15). Professor: Brystkræftscreening med AI giver hurtigere resultater, men det øger risikoen for fejl. *Altinget*. <https://www.alinget.dk/digital/artikel/professor-brystkraeftscreening-med-ai-giver-hurtigere-resultater-men-det-oeger-risikoen-for-fejl>
- Latour, B. (1999). *Pandora's hope: Essays on the reality of science studies*. Harvard University Press.
- Lauritzen, A. D., Rodríguez-Ruiz, A., von Euler-Chelpin, M. C., Lynge, E., Vejborg, I., Nielsen, M., & Lillholm, M. (2022). An artificial intelligence-based mammography screening protocol for breast cancer: Outcome and radiologist workload. *Radiology*, 304(1), 41–49. <https://doi.org/10.1148/radiol.210948>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>
- Perez, C. C. (2019). *Invisible women: Data bias in a world designed for men*. Abrams Press.
- Regeringen, KL og Danske Regioner. (2025). *Mere tid til det vigtige. Målbillede for udrulning og anvendelse af kunstig intelligens i den offentlige sektor*.
- Rosen, A. W., Ose, I., Gögenur, M., Andersen, L. P. K., Bojesen, R. D., Vogelsang, R. P., & Gögenur, I. (2025). Clinical implementation of an AI-based prediction model for decision support for patients undergoing colorectal cancer surgery. *Nature Medicine*, 31(11), 3737–3748. <https://doi.org/10.1038/s41591-025-03612-4>
- Suchman, L. (2023). The uncontroversial 'thingness' of AI. *Big Data & Society*, 10(2), 1–5. <https://doi.org/10.1177/20539517231206794>

- Sundhedsdatastyrelsen. (2018). *Strategi for digital sundhed 2018–2022*. <https://sundhedsdatastyrelsen.dk/da/strategier-og-projekter/strategi-for-digital-sundhed>
- Svensson, A. M., & Jotterand, F. (2022). Doctor ex machina: A critical assessment of the use of artificial intelligence in health care. *The Journal of Medicine and Philosophy*, 47(1), 155–178. <https://doi.org/10.1093/jmp/jhab036>
- Tupasela, A., Snell, K., & Tarkkala, H. (2020). The Nordic data imaginary. *Big Data & Society*, 7(1), 1–11. <https://doi.org/10.1177/2053951720907107>
- Vetter, D., Amann, J., Bruneault, F., et al. (2023). Lessons learned from assessing trustworthy AI in practice. *Data and Information Science*, 2, 35. <https://doi.org/10.1007/s44206-023-00063-1>
- Vikkelsø, S. (2005). Subtle redistribution of work, attention and risks: Electronic patient records and organisational consequences. *Scandinavian Journal of Information Systems*, 17(1), 3–30.
- Wen, D., Soltan, A., Trucco, E., & Matin, R. N. (2024). From data to diagnosis: Skin cancer image datasets for artificial intelligence. *Clinical and Experimental Dermatology*, 49(7), 675–685. <https://doi.org/10.1093/ced/llae112>
- Zajac, H. D., Li, D., Dai, X., Carlsen, J. F., Kensing, F., & Andersen, T. O. (2023). Clinician-facing AI in the wild: Taking stock of the sociotechnical challenges and opportunities for HCI. *ACM Transactions on Computer-Human Interaction*, 30(2), 1–39. <https://doi.org/10.1145/3571810>

# Before Deployment

## Anticipatory Infrastructuring and Early AI Integration in Acute Stroke-ready MRI Workflow

Maria Bach Nielsen<sup>1, 2</sup>

Tom Børsen<sup>1</sup>

Casper Knudsen<sup>3</sup>

---

<sup>1</sup>Responsible Technology Futures, Department of Sustainability & Planning, Aalborg University

<sup>2</sup>Cerebriu, Copenhagen

<sup>3</sup>Techno-Anthropology and Participation, Department of Sustainability & Planning, Aalborg University

*mariabachnielsen@plan.aau.dk*

*boersen@plan.aau.dk*

*casper@plan.aau.dk*

Nielsen, Maria Bach; Børsen, Tom & Knudsen, Casper. 2026. Before Deployment: Anticipatory Infrastructuring and Early AI Integration in Acute Stroke-ready MRI Workflow. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 19-47. DOI: 10.7146/tfss.v25i44.156503

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

*Artificial intelligence (AI) decision support is often framed as a technical response to diagnostic pressure, yet making AI usable in clinical settings depends on anticipatory work that connects new outputs to existing infrastructures, routines, and responsibility relations. This article examines early integration work around Apollo, a commercially available AI-based decision-support system for brain MRI, explored in a Danish acute stroke-ready*

*hospital context as support for triage and attention during image acquisition rather than as diagnostic authority. The analysis draws on longitudinal rapid ethnographic fieldwork (2019–2024), including approximately 60 hours of observation across collaboration meetings, clinical workflow settings, training, and pilot-related sessions; six semi-structured interviews with key stakeholders; and a corpus of project documents. Focusing on pilot and pre-implementation activity, the article traces how clinicians, developers, and institutional actors negotiated what Apollo could be allowed to do, how it could connect to the installed base, and how its outputs could be made visible and timely within established work practices. Conceptually, the article frames these activities as anticipatory infrastructuring: alignment work oriented toward possible future routine use, carried out while attachments and responsibilities remain unsettled. The article contributes an empirically grounded account of anticipatory infrastructuring in early clinical AI integration. It shows how, in pilot and pre-implementation work, actors collectively stabilized a future possibility of use by (1) delimiting scope and legitimacy, (2) stabilizing routability and tempo of data flows, (3) tuning visibility within existing screen ecologies, and (4) negotiating expectations for noticing and responding to algorithmic suggestions.*

## Før udrulning: Anticipatorisk infrastrukturering og tidlig AI-integration i akut apopleksi MR-arbejdsgang

*Kunstig intelligens (AI) til beslutningsstøtte fremstilles ofte som et teknisk svar på diagnostisk pres, men at gøre AI anvendelig i kliniske sammenhænge kræver anticipatorisk arbejde, hvor nye resultater kobles til eksisterende infrastrukturer, rutiner og ansvarskæder. Denne artikel undersøger tidligt integrationsarbejde omkring Apollo, et kommercielt tilgængeligt AI-baseret beslutningsstøttesystem til hjerne-MR, som på et dansk hospital med apopleksiafsnit uden trombolyse- eller trombektomimulighed blev udforsket som støtte til triagering og fokusering under billedoptagelse snarere end som diagnostisk autoritet. Analysen bygger på longitudinelt 'rapid' etnografisk feltarbejde (2019–2024), herunder cirka 60 timers observation på tværs af samarbejds møder, kliniske arbejdsgangssituationer, træning og pilotrelaterede sessioner; seks semistrukturerede interviews med centrale aktører; samt et dokumentkorpus af projektmaterialer. Med fokus på pilot- og præimplementeringsaktivitet følger artiklen, hvordan klinikere, udviklere og organisatoriske aktører forhandlede, hvad Apollo kunne få lov til at gøre, hvordan systemet kunne kobles til den etablerede tekniske infrastruktur, og hvordan systemets resultater kunne gøres synlige og rettidige i etablerede arbejdsgange. Analytisk forstås dette som 'anticipatorisk infrastrukturering': et tilpasningsarbejde rettet mod mulig fremtidig rutinebrug, gennemført mens forbindelser og ansvar stadig er uafklarede. Artiklen*

*bidrager med en empirisk funderet fremstilling af 'anticipatorisk infrastrukturering' i tidlig integration af klinisk AI. Den viser, hvordan aktører i pilot- og præimplementeringsarbejde kollektivt stabiliserede en fremtidig mulighed for brug ved (1) at afgrænse systemets omfang og legitimitet, (2) at stabilisere dataflowenes tempo og vej gennem systemet, (3) at justere synlighed inden for eksisterende skærmøkologier og (4) at forhandle forventninger til at opdage og reagere på algoritmiske forslag.*

## Introduction

Clinical AI is often promoted as a response to diagnostic pressure, yet turning an AI model into usable decision support depends on how its outputs can be connected to existing infrastructures, routines, and the allocation of responsibility and accountability for acting on them. Recent reviews emphasise workflow feasibility and integration as recurring challenges, while noting limited empirical detail on the practical work that makes AI workable in situ. Khan et al (2024) identified workflow feasibility and integration into existing clinical workflows as recurring concerns in AI implementation, while noting that the practical work of embedding tools into everyday routines is less consistently addressed. Peek et al. (2024) similarly argued that implementation-oriented medical informatics research on AI-based clinical decision support remains fragmented and often technology-centric with limited empirical detail on the organizational and infrastructural changes involved in successful deployment. This article examines early integration work around Apollo, a commercially available AI-based decision-support system for brain MRI, explored in a Danish acute stroke-ready setting as support for triage and attention during image acquisition rather than as diagnostic authority. We follow Apollo through feasibility, pilot, and pre-implementation activity (2019–2024), before any stabilised routine use.

We conceptualise this early integration as anticipatory infrastructuring, i.e., distributed alignment work oriented toward possible future routine use, carried out while points of attachment, visibility, timing, and responsibilities remain unsettled. The article asks: What role does anticipatory infrastructuring play in the early integration of an MRI decision-support system, and with what implications for accountability, workflow fit, and trust-in-use? We show how participants stabilised a possible future use through four coupled moves: (1) delimiting scope and legitimacy, (2) stabilising routability and tempo of data flows, (3) tuning visibility within existing screen ecologies, and (4) negotiating expectations for noticing, interpreting, and responding to algorithmic suggestions.

## Related work

Infrastructures are relational achievements embedded in practice and learned as part of membership, becoming visible primarily in moments of breakdown or when new elements must connect to an existing installed base, such as standards, access arrangements, or information system (Star, 1999; Star & Ruhleder, 1996). This perspective foregrounds integration work as a practical and social accomplishment rather than a purely technical task.

Edwards et al. (2007) emphasized that infrastructure is rarely built from scratch; it develops historically and must be extended through existing organizational and technical arrangements rather than implemented as a clean, blueprint design. Building on this orientation, Pipek and Wulf (2009) propose *infrastructuring* to foreground the ongoing work of developing, adapting, and maintaining infrastructure over time. In healthcare settings, this work is often distributed across heterogeneous actors and institutional boundaries. For example, Aanestad et al. (2017) show how eHealth infrastructures are shaped through negotiation and incremental alignment, while Bossen and Markussen (2010) demonstrate how breakdowns and upgrades in medication infrastructures make visible the extensive coordination work and ordering of devices through which care and accountability are materially organized.

These accounts provide the context for the analysis of AI integration in imaging as ongoing alignment work that connects new outputs to existing coordination routines, documentation practices, and responsibility relations. Sociotechnical frameworks for healthcare AI emphasize that clinical impact depends not on model performance alone but on how AI becomes part of a broader intervention ensemble embedded in workflows, organizational arrangements, and governance practices. McCradden et al. (2023) make this point explicitly by framing clinical AI as inseparable from the surrounding practices and institutional conditions that enable it to function safely and meaningfully. We build on this orientation by analysing early AI integration as anticipatory infrastructuring. Methodologically, Karasti et al. (2018) also highlight that infrastructures take time and are often empirically knowable through intervention, by following how connections are assembled across settings, standards, and routines. These accounts point to a methodological implication for clinical AI. Early integration involves practical decisions about what counts as sufficiently interpretable to act on, and about how accountability is distributed when system reasoning is only partly inspectable.

Introducing AI into clinical decision-making, therefore, raises questions not only about performance, but also about how clinicians can account for decisions

partly shaped by algorithmic output. In a study of diagnostic radiologists using AI, Lebovitz et al. (2022) showed that clinicians did not simply trust or reject AI systems; instead, they developed AI interrogation practices to relate algorithmic outputs to professional judgement. Access to such practices was uneven, shaping whether AI became meaningfully engaged in diagnostic work. Fox et al. (2023) similarly showed that integration in work relies on patchwork labour, ongoing adjustment, coordination, and maintenance that bridge the gap between what systems promise and what they can actually do in practice. However, comparatively little empirical work traces what happens before an AI system becomes routine clinical infrastructure, i.e., what negotiations, constraints, and alignment work through which a system's role, scope, and points of connection are defined in pilot and pre-implementation settings.

Bansler and Havn (2010) argue that pilots are not small-scale versions of full implementations but rather distinctive socio-technical arrangements in which scope, organizational commitment, and learning become central challenges. Building on this orientation, Malm-Nicolaisen et al. (2024) document how successful integration outcomes can conceal extensive preparatory work, including negotiation, technical adaptation, and cross-stakeholder coordination prior to routine use. In radiology specifically, Farič et al. (2024) provide early accounts of integrating AI decision support into clinical practice, while Ramsay et al. (2025) analyse procurement and early deployment as socio-technical work rather than merely an administrative prelude. Finally, Vasey et al. (2022) emphasize that early-stage clinical evaluation should account for human factors, workflow realities, and safety concerns in situ. These bodies of work point to a clear need for empirical studies that follow the early integration work through which AI is made potentially usable and accountable, before use stabilizes as a routine, taken-for-granted practice.

Across adjacent literatures, the work of anticipation has been in focus, naming the skilled, distributed practices through which actors make futures actionable in the present. Steinhardt and Jackson (2015) foreground anticipation work as cultivating and channelling expectations, designing pathways into imagined futures, and maintaining those visions under changing conditions. In high-reliability settings, Johansen et al. (2016) has analysed anticipation more explicitly as proactive 'what can possibly go wrong' work, showing how engineers, researchers and technicians anticipate and mitigate possible problems and how the nature of anticipatory work shifts between planning and operational phases. Nielsen et al. (2022) further show that the ongoing narrative labour of adapting forward-looking narratives in response to delays, conflict, and setbacks is crucial to public-

sector digital transformation. In healthcare datafication, Helén and Tarkkala (2024) propose that anticipation in data-driven healthcare is organized through practical devices that make expectations actionable. This article uses these perspectives to define *anticipatory infrastructuring* in early AI integration. Specifically, as pilots and pre-implementation are not merely preparatory stages but arenas where actors collectively stabilize particular futures by making breakdowns, mismatches, and responsibility dilemmas visible early, and by converting visions of AI-enabled care into concrete socio-technical attachments, workflows, interfaces, data routes, and accountability arrangements through which the future infrastructure of professional practice (including its divisions of labour and (in)visibilities) is actively negotiated before routine use.

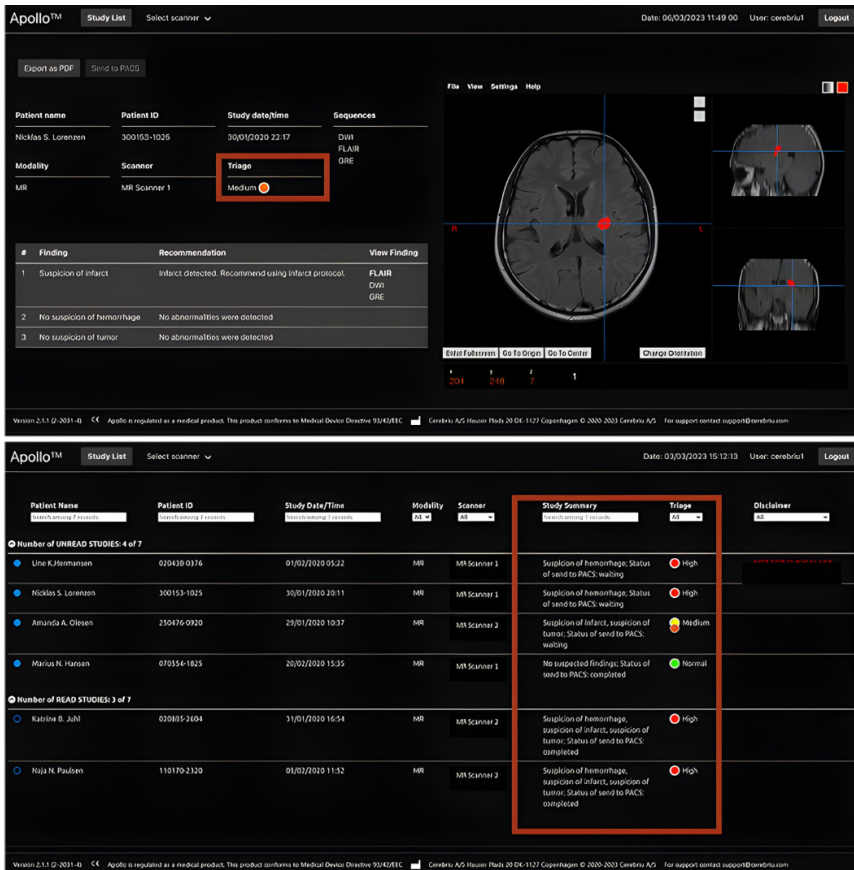
We use the term anticipatory infrastructuring to describe early-stage alignment work conducted in pilots and pre-implementation settings where the future use of a system remains open. Pilots are not small-scale versions of full implementations but rather distinctive socio-technical arrangements in which scope, commitment, and organisational learning are actively negotiated. Analytically, we attend to how actors delimit what Apollo may legitimately influence, stabilise routable and timely data flows, tune visibility within existing screen ecologies, and negotiate responsibilities for noticing, interpreting, and responding to algorithmic suggestions.

## Introducing the case of Apollo

### *Acute stroke-ready MRI as an infrastructural setting*

The case is situated in a Danish hospital environment that, from 2019 onward, was locally described as moving toward an MRI-first approach for acute neurology, reorganising pathways while still relying on other modalities. Imaging work is distributed across scanners, Picture Archiving and Communication Systems (PACS), Radiological Information Systems (RIS), Electronic Health Records (EHR), clinical handovers, and time-sensitive triage routines. Introducing AI here is therefore not simply a matter of inserting a new decision-support output; it requires negotiating how AI output aligns with temporal rhythms, visual work practices, and chains of responsibility, especially when outputs appear during acquisition work where speed and coordination matter. Apollo is a commercially available AI-based decision-support system for brain MRI. It analyses sequences

such as diffusion-weighted imaging (DWI), T2-weighted Fluid-Attenuated Inversion Recovery (T2 FLAIR), and susceptibility-weighted imaging (SWI) and produces output during acquisition to support triage and attention under clinician interpretation (see Figure 1). In validation work, Krag et al. (2023) evaluated Apollo as a CE-marked tool intended to support detection of time-critical findings on brain MRI, including acute ischemic lesions.



*Figure 1: Highlighted View of Apollo Version 2.1.2 Triage and Smart Backlog. Patient data shown is simulated for demonstration purposes<sup>1</sup>*

Apollo is analytically useful here because its intended role is situated at the point of image acquisition and prioritization, where coordination, timing, and

1. Figure 1, 4 and 8 and table 1 and 2 are created by first author in project management capacity at Cerebriu A/S (2020-2023).

responsibilities are tightly coupled across radiographers, radiologists, and downstream clinical actors. This positioning brings early integration work into view as questions about when outputs should appear, who should notice them, and how they should be handled alongside existing routines and documentation practices. Accordingly, early work centered on workflow fit, timing, and responsibility before routine use was established.

## Methods and analytical lens

This article reports a longitudinal qualitative case study (2019–2024) of early integration work around Apollo. The first author conducted multi-sited rapid ethnographic fieldwork across collaboration meetings, pilot-related sessions, and clinical workflow settings, complemented by interviews, workshops, and document analysis. Rapid ethnography is well suited to time-pressured clinical environments because it enables focused attention to situated coordination and breakdowns while still producing a longitudinal account (Vindrola-Padros & Vindrola-Padros, 2018).

The empirical corpus combines participant observation, interviews, workshops, and document analysis across the Apollo collaboration (2019–2024). Between 2019 and 2021, the first author conducted approximately 13 hours of participant observation in company-hospital collaboration settings, including project management meetings, workflow design discussions, IT and governance alignment meetings, and simulated live-test sessions. This period also involved close engagement with internal company documents and planning materials and coincided with early feasibility work and scoping of the collaboration.

From 2021 to 2023, data collection shifted toward rapid ethnographic observation of clinical workflow, comprising approximately 45 hours of in-situ observations in neurology and radiology departments. These observations focused on routine coordination, documentation practices, and temporal organization of stroke-adjacent imaging work, as well as brief observations during Apollo test phases and pilot-related technical feedback loops. Follow-up observations in 2023 captured how infrastructural reconfigurations and organizational changes redistributed coordination work over time.

In parallel, six semi-structured interviews were conducted in 2021 with key stakeholders, including radiographers, radiologists, neurologists, and innovation staff. Each interview lasted approximately 30 minutes and was recorded and transcribed verbatim. Interviews were used to complement observational material

by eliciting participants' accounts of responsibility, practical constraints, and perceived implications of early AI integration. From 2022 to 2024, the empirical material was further extended through participation in three 60–90-minute reflective workshops and staging sessions (including pilot debriefs and planning sessions), involving radiologists, radiographers, neurologists, the hospital management and their innovation department, and company representatives. These workshops provided occasions to revisit observations, rehearse integration scenarios, and negotiate scope, accountability, and workflow fit.

Across the entire period, a document corpus of more than 30 items was collected and analysed, including internal reports, presentations, pilot plans, meeting summaries, and legal and regulatory materials. These documents were used to trace how integration problems, responsibilities, and proposed solutions were articulated and re-articulated over time. Fieldnotes were written contemporaneously and expanded immediately after each engagement. All materials were handled in de-identified form in the research corpus, and the analysis reported here does not include patient-identifying data.

The first author conducted the primary fieldwork from a position of embedded engagement (industrial PhD/project management), enabling access across clinical and development contexts but also requiring reflexive attention to how researcher involvement shaped interactions and interpretation. She addressed this through systematic field note practices, triangulation across observations, interviews, and documents, and iterative analytic memoing. Analysis proceeded iteratively. Fieldnotes, interview transcripts, and documents were open coded for episodes where Apollo's introduction generated negotiation, adjustment, or breakdown related to (a) workflow fit and timing, (b) visibility and interpretability of outputs, (c) allocation of responsibility for noticing and acting, and (d) infrastructural connections and data pathways. Codes were compared across actor groups and settings and refined into higher-level themes capturing recurring forms of early integration work which structure the following sections.

## Findings: Early integration work

In the earliest discussions around Apollo in 2019, the system was not yet stabilized as an MRI tool in any concrete sense. Early negotiations centered on how AI could meaningfully enter MRI workflows without requiring the hospital to reorganize clinical responsibility around the system.



*Figure 2: Observation of early discussions with hospital staff about workflow at the scanner post proof-of-concept meeting<sup>2</sup>*

The section is organized around four coupled moves through which participants worked to stabilize the socio-technical conditions for possible future use.

### *From broad AI promise to a workable clinical niche*

This narrowing began already in initial feasibility dialogues and continued through a live test design workshop (Sep 2020), where clinicians, innovation staff, and developers drafted an initial workflow and data-sharing logic and articulated where AI output might plausibly appear and matter. Subsequent feedback sessions further narrowed the focus, explicitly framing ethical feasibility and non interventional positioning as necessary conditions for moving forward.

These workshops and feedback session played an important foundational role in the anticipatory alignment between actors exploring what kind of intervention was

---

2. Figure 2, 3, 5, 6 and 7 are images based on fieldwork photographs and hand-drawn sketches, transformed using OpenAI's ChatGPT (2024) to create anonymized, sketch-style illustrations. Original materials retained in research corpus.

being attempted. Before infrastructure was built, the collaboration had to decide what integration would even mean: a triage cue? A quality check? Something that changes protocols? Something that only creates a quicker overview?



*Figure 3: Illustration of observed meeting with innovation unit*

This work form primarily concerns accountability and scope. How Apollo could support workflow ordering without becoming a source of diagnostic authority, and how responsibilities remain intelligible across roles. Throughout this early work, Apollo’s legitimacy was repeatedly tied to maintaining professional judgment and avoiding any impression that responsibility was being shifted to the system. This became a central tension shaping what could be proposed in pilots, shown in interfaces, and claimed in evaluation. Interview and observation material captures this boundary work directly. One radiologist summarized Apollo as “an extra pair of eyes, but it shouldn’t speak for us” (Radiologist, Fieldnotes, 2022). A neurologist made the same distinction: “We need something that supports decision-making, not something that decides” (Neurologist, Fieldnotes, 2022). When discussions moved into clinical spaces, interpretability became central. During a training session, radiologists asked how they would know why a scan is flagged

as urgent. When told that the system did not yet display its rationale, one clinician replied: *"If I can't see what the AI sees, how do I trust it? I still have to go through the scan myself"* (Radiologist, Fieldnotes from Training Session, 2020). Across meetings, these statements functioned as constraints that shaped early design talk. Support could be acceptable if it did not generate new obligations that were impossible to meet in an acute workflow, e.g., if it did not require clinicians to defend an AI reading as if it were a diagnosis or introduce a new step that slowed down the scan. Conversely, support became suspicious when participants imagined it as an alert that someone must always verify, or as an output that could be used retrospectively to blame clinicians for not acting on algorithmic suggestions. These worries made accountability an early integration question long before any routine use was at stake.

A third and persistent strand of early scoping involved making Apollo into an object that could be governed within Danish and European legal/regulatory frameworks and within hospital accountability structures. In meetings with hospital leadership, Apollo was often framed less as a tool and more as a governance object. Something that required clarity on responsibility, compliance, and institutional risk before anyone could even speak meaningfully about implementation. This is captured in an administrator's statement during an early proof of concept meeting: *"Before we even talk about implementation, we need clarity on how this fits within Danish and European regulatory frameworks. It's not just about whether AI works, but whether we can legally and ethically use it in patient care."* (Head of Radiology department, Early proof of concept meeting, Fieldnotes, 2020). Within the collaboration, participants also had to negotiate what would count as meaningful progress. These negotiations were not only between clinic and company, but also within the hospital (innovation, clinicians, IT/legal), where actors held different stakes in what could reasonably be promised and measured.

Through a series of meetings, the collaboration experimented with evaluation vocabularies. A KPI Brainstorm held in November 2020 was particularly telling because it marks an explicit shift away from purely quantitative measures and toward workflow evaluation as a credible way of assessing value in an acute stroke-ready MRI context (see Figure 4).

Empirically, this shift matters because it reframes the system. Apollo becomes less an object to be assessed primarily by model performance and more a socio-technical proposition whose viability depends on whether it can be aligned with existing tempo, coordination, and clinical judgment practices. This also helps explain why the early integration work did not converge on a single definition of what Apollo is. A system can simultaneously be pitched as efficiency support

and treated as a governance risk; it can be imagined as triage support while being questioned as a potential administrative burden. Early integration work, in this sense, included negotiating which of these framings is allowed to matter in decision-making about the project.

*Figure 4: Outcome slide from Apollo KPI brainstorm, 26th Nov 2020, with a minor annotation for anonymization (Cerebriu A/S, 2020)*

## KPIs & Metrics:

	Qualitative	Quantitative
Radiology	<ul style="list-style-type: none"> <li>Workflow improvement of acute patients (neuroradiologist triages directly w/o radiographer intervention)</li> <li>Better (less stressful) review of neuroradiologist's Monday morning backlog (Triage)</li> </ul>	<ul style="list-style-type: none"> <li>Reduced number of unnecessary admissions (KPI: Number of cases of no-stroke-findings following admission of TIA during nightshift as confirmed by next-day MRI)</li> <li>Number of scans saved (tumour protocol aborted)</li> <li>Number of Gd doses saved (tumour protocol aborted)</li> <li>Faster review of neuroradiologist's case load (triaging)</li> </ul>
Neurology	<ul style="list-style-type: none"> <li>Improved confidence in decisions whilst triaging</li> <li>Example article: "Confidence improvement of junior clinicians' decision-making during night shifts when discharging vs admitting asymptomatic TIA patients and using Apollo as a proxy for neuroradiological reporting"</li> </ul>	<ul style="list-style-type: none"> <li>Number of patients triaged successfully before 10pm cutoff for admission to neurodept</li> <li>Number of patients treated correctly first time (TIA vs stroke)</li> <li>Time-to-medication-start (blood thinners)</li> </ul>



Another practical achievement of this early scoping work was the production of representations that could travel across actors and settings. Workflow maps, pilot plans, and slide decks were not neutral documentation; they were working objects that helped align what clinicians meant by workflow, what developers meant by integration, and what administrators meant by responsibility and compliance.

The Live Test Design workshop (Sep 2020) and subsequent feedback sessions produced the first workflow maps of stroke imaging practices. Later, Neurology Workflow Review (Nov 2020) and the Workflow Mapping Interviews (May–Jun 2021) further supported this representational work by identifying interdepartmental coordination challenges and clarifying where and how an AI tool could plausibly fit into existing routines. This work is visible in early technical alignment discussions involving hospital innovation units, regional IT, and development teams. The Alignment Meeting (Hospital IT Company, Jan 2021)

outlined PACS connectivity and server setup; the Configuration Meeting (Jan 2021) decided on automatic sequence transfer (auto-push); and subsequent technical infrastructure meetings surfaced SSL and access rights as barriers that had to be resolved. Importantly, these were not treated as mere technicalities. They were discussed as conditions for responsible conduct and institutional accountability, i.e., who can access what, under what permissions, and whether transfer paths can be justified and audited. In this sense, early infrastructural planning was part of scoping; it contributed to defining what Apollo could become (and what it could not become) within the hospital's installed base.

By the end of 2021, the collaboration had produced the legal, infrastructural, and relational preconditions that allowed the first on site pilot and user training to begin in early 2022. This included: (a) a non interventional framing that could be defended institutionally, (b) governance milestones enabling retrospective evaluation and continued planning, and (c) technical and organizational agreements about how data should flow and how the project should be coordinated. These achievements did not settle the central practical questions that would later dominate situated integration work, where outputs should appear, who should see them, how quickly they should surface, and what it would mean, in accountability terms, to notice or not notice an algorithmic suggestion. Those questions were not waiting for the pilot. They were already being shaped here, but they became materially concrete only once attention turned from legitimacy and scope toward workflow constraints and infrastructural attachment in practice.

*Workflow constraints as early integration work: mapping coordination, temporalities, and infrastructural joints*

Early integration work did not only consist of meetings about Apollo. Because Apollo's proposed role was to provide real-time support during image acquisition and triage, integration also required making the constraints of their stroke MRI workflow empirically visible. How urgency is produced and negotiated, where coordination breaks down, and infrastructural connections determine whether information arrives in time to matter. Workflow constraints therefore functioned as anticipatory infrastructuring in its own right. It became a form of pre-deployment alignment work that clarified which outputs could become timely, visible, and responsibly actionable within an already functioning clinical ecology. Across fieldwork in emergency neurology and radiology (observations in 2021 and follow-up in 2023), workflow appeared as a distributed choreography, sustained through continuous articulation and adjustment across people, spaces, and

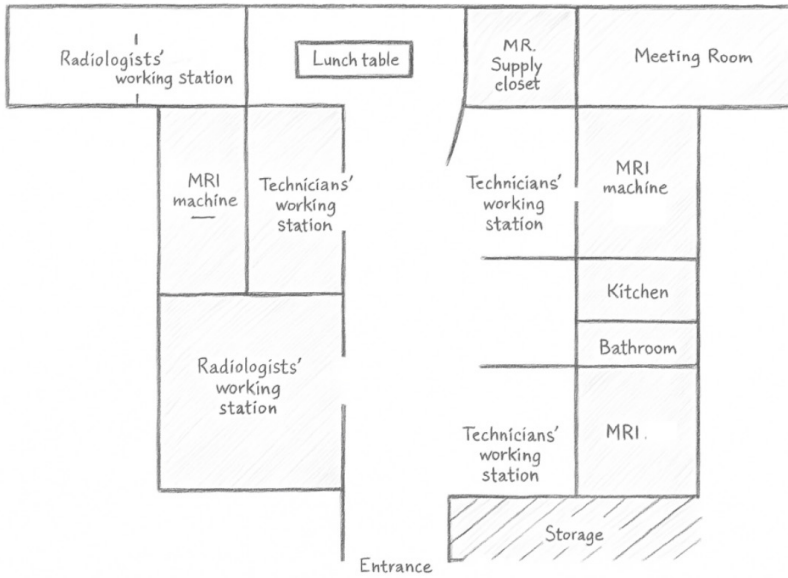


Figure 5: Sketch of radiology department




Figure 6: The Radiographer Workstation

infrastructures. This is included because integration depended on whether outputs could attach to these rhythms and junctions, not only on interface or model design.

A central finding from workflow observations is the extent to which radiographers mediate between clinical urgency and technical/organizational feasibility. Inside the MRI suite (see figure 6), scan acquisition was not a self-contained technical process. It was a collaborative situation where neurologists, radiographers, and radiologists negotiated what needed to be seen and how quickly.

These situations also show that urgency is not a stable label that simply travels through a booking system. It has to be translated into the practicalities of scanner time, patient mobility, staffing, and queue management. Radiographers described how competing temporalities collide in everyday work. Fieldnotes captured phone-driven escalation and triage conflicts: *“Neurology just called again, they want their stroke patient now, but the trauma case is already waiting.”* (Radiographer, Workflow observation in ER, Fieldnotes, 2022). Another case involved a stroke-suspected patient mislabelled in the booking system and slipping down the queue, only to be escalated by phone when a neurologist intervened. In interview form, this labour was articulated explicitly as translation work: *“We translate what ‘urgent’ means, everyone says their patient is urgent, but the timing depends on the scanner, on contrast, on who can move the patient.”* (Radiographer, Workflow observation in Radiology Unit, Fieldnotes, 2022).

For Apollo, this meant that an urgent AI signal could only function as a negotiable cue because urgency was already continuously translated and contested through phone calls, queue management, staffing constraints, and competing clinical priorities. This matters because any proposal to deliver real-time AI output during acquisition enters a workflow where urgency is already continuously negotiated. Integration in this sense is not adding one more signal; it is fitting AI output into an existing attention economy and temporal choreography without creating new obligations that staff cannot meet (e.g., additional checks, new escalation steps, or ambiguous responsibility for acting on alerts). This constraint helps explain why later pilot and planning work repeatedly returned to boundary-making around what Apollo could legitimately demand of attention and response, and to tempo alignment as a condition for usefulness.

Workflow constraints also made visible that coordination work in stroke care is often dominated by documentation, locating patients/information, and filtering signals. In emergency neurology, the workstation appeared as a compact coordination hub organized by roles and yet constantly in motion. Clinicians moved between documentation, phone calls, brief consultations across desks, and

patient encounters. The rhythm of the day was anticipatory and reactive, mornings used to get ahead in documentation, with midday arrivals pulling attention into real-time triage. A telling episode concerned a suspected transient ischaemic attack (TIA) case where the patient's location had not been logged correctly in the digital system. Despite repeated calls, the on-call consultant (N-BV) could not locate the patient through the system and eventually searched examination rooms in person. The delay was not caused by a lack of clinical competence or a lack of imaging technology, but by fractured coordination across digital and physical spaces, which increased the documentation burden. Infrastructural promise (visibility and speed) became infrastructural friction (uncertainty and delay), and workarounds (phone calls, physical searching) carried the case forward.

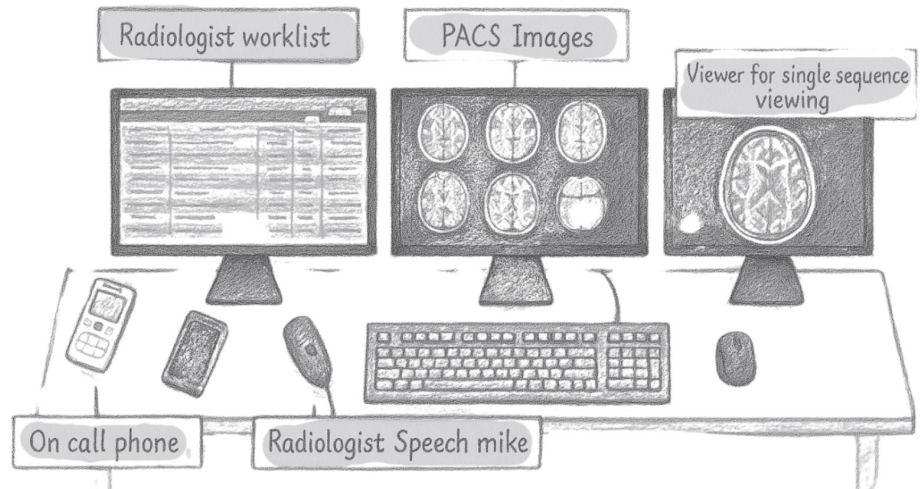
For Apollo, this means AI output would enter an already-filtered attention ecology, so its visibility and formatting would need to avoid becoming background noise while also not creating a new category of alerts that staff feel compelled to verify regardless of workload.

Adding an AI output also changes what kind of signal it becomes in an environment already saturated with alerts, documentation requirements, and coordination breakdowns. The question becomes: Will AI output be filtered out like background noise, or will it create new responsibility obligations that clinicians and radiographers cannot absorb?

Across both neurology and radiology, breakdowns repeatedly appeared not inside single devices, but at the interfaces between systems, roles, and locations. This was expressed bluntly by an IT technician: *"It's never the scan itself that fails, it's always the connection."* (Local hospital IT technician, workflow observation in ER, Fieldnotes, 2022). Freezing PACS interfaces, misaligned bookings, and delayed image transfers disrupted the choreography of triage and interpretation. These frictions also foreshadowed a practical integration problem for Apollo: if AI depends on stable, timely transfer of sequences to an AI server and then back into clinical systems, then connectivity becomes a primary feasibility condition, because delays push outputs outside the clinical window where triage decisions are still live. In this sense, the installed base did not function as a neutral delivery channel for AI output; it was the site where early integration either succeeded as routable and timely support or failed as an unreliable side channel.

Remote diagnostics (accelerated during the COVID-19 period) introduced additional coordination frictions. Radiographers on site sometimes needed rapid clarification when artifacts appeared or when unexpected scan quality issues emerged, but the radiologist was not physically present. Although messaging was technically possible, it was experienced as insufficient for showing and pointing

in real time. A radiographer described this limitation: “Sometimes you just want to show them the image, point to the artifact, and ask, ‘What do we do now? But they’re not here’” (Radiographer, Rapid workflow observation in Radiology Unit, Fieldnotes, 2021).



*Figure 7: Radiologist working station*

This set of dependencies is directly relevant to Apollo’s early integration because it clarifies what workflow fit entails materially. Not just user acceptance, but timing, routing, and reliability across PACS/RIS/EHR and across distributed clinical locations. It also helps explain why later training and planning discussions repeatedly returned to questions such as how outputs would appear in PACS and who would be responsible for verifying alerts when time is short. For Apollo, the 2023 reconfiguration shows that infrastructural change redistributes rather than removes attachment work, meaning an AI system would need to remain workable across shifting coordination points (booking interfaces, porter capacity, bedside logistics, and control-room negotiations) rather than assuming one stable workflow to integrate into. These observations show why early Apollo integration could not be addressed solely through interface design or governance decisions in meetings. The workflow constraints of stroke-adjacent MRI, negotiated urgency, limited attention capacity, and fragile infrastructural junctions shaped what kinds of AI outputs could become timely, visible, and responsibly actionable in practice. The next section shows how these constraints were brought into reflective planning and pilot activity, where clinicians, developers, and institutional actors collectively

explored how Apollo could attach to existing routines without becoming a new source of interruption, uncertainty, or responsibility burden.

*Reflective planning as early integration work: co-creating conditions for use*

Building on the workflow constraints described above, early integration also required creating spaces where clinicians, developers, and institutional actors could translate field insights into concrete decisions about where Apollo could attach, what it should make visible, and how responsibility should remain organized. In the Apollo case, reflective planning did not finalize implementation. Instead, it functioned as anticipatory infrastructuring: collective alignment work oriented toward possible future routine use, conducted while the system’s role, timing, and points of connection were still unsettled.




Figure 8: Pilot observation setup (Cerebrii A/S)

This work form primarily concerns trust-in-use and workflow fit: tuning when and how outputs become visible so that they can be consulted without disrupting established screen ecologies and time-critical coordination. Early groundwork for reflective planning emerged during preparation for a two-day internal pilot at the university hospital in January 2022. The pilot was explicitly exploratory – the focus was not diagnostic performance, but whether Apollo could operate under clinical conditions and how professionals would relate its outputs to their everyday work. Preparation began with a user training session in mid-January 2022, preceded by a short presentation that framed Apollo as supporting prioritization rather than replacing clinical judgment.

The first hands-on session took place in the radiology department, set up as an on-site demonstration at the clinicians’ workstation. Participants included a neuroradiologist, a clinical researcher, and a hospital project coordinator from the innovation unit. The system was accessed via its dedicated server and configured to receive MRI scans automatically via an auto push setup. To make Apollo visible without displacing the existing diagnostic setup, an additional laptop screen was connected, producing a four-screen configuration that mirrored the department’s established screen ecology (see Table 1 below).

*Table 1: Configuration on screens at radiology department*

<b>Screen 1:</b>	Patient lists and administrative systems (such as PACS and RIS, used for storing images and managing reports)
<b>Screen 2:</b>	The main screen for reading MRI images
<b>Screen 3:</b>	Comparison images or prior patient studies
<b>Screen 4</b>	(Added): The pilot AI triage system interface

What mattered in this session was how the system entered the situation: where it sat, what it demanded of attention, and whether it could be consulted without re-engineering overview.

Very early, the session turned into an infrastructuring problem. Eleven patient cases were sent to the system for processing, but none completed analysis within the first hour. Investigation showed that the Diffusion Weighted Imaging (DWI)

sequence, critical for stroke, was not transferred correctly when scans were manually uploaded. The problem disappeared when switching to auto push, and this became a practical conclusion: Future pilots should rely on automated transfer until manual upload could be debugged. This episode is analytically important because it shows that early integration hinges on aligning system tempo with clinical tempo. A system can be validated in a technical sense and still fail as a clinical proposition if the data pathway introduces delay, uncertainty, or friction at the wrong moment. Here, integration work became about making the connection reliable enough that Apollo’s outputs could be taken seriously as part of the work, rather than dismissed as not ready because results appear too late or not at all.

Once the transfer problem was resolved, discussion quickly moved to how Apollo should appear in the radiologist’s working environment. The neuroradiologist emphasized that an additional screen for the AI interface was preferable to opening browser tabs inside existing screens precisely because the established visual layout supports diagnostic concentration and overview. In this setting, integration was not simply about providing an output, but about whether the output could be consulted without interrupting the visual and cognitive organization of the workstation. This attention to small interface details surfaced as practical concerns about overview, clutter, and cognitive load. The neuroradiologist suggested concrete adjustments (Table 2 below).

*Table 2: Improvement suggestions (Cerebriu A/S)*

#	Improvement suggestions from the neuroradiologist
1	The system should allow users to mark cases as “in progress” or “completed” to maintain a clear overview and reduce visual clutter.
2	Cases already reviewed should automatically disappear from the active list.
3	Sorting preferences (such as by scan date or time) should be remembered between sessions.
4	Colour overlays used to indicate priority might be replaced with directional arrows to reduce visual stress.

These suggestions were not only usability preferences. They articulated how trust and attention are sustained in practice: by keeping the interface legible and by preventing the AI list from becoming a second queue that clinicians must babysit.

A second internal familiarization session one week later built on these adjustments. With the DWI issue resolved and auto-push functioning consistently, clinicians experimented with the question of when Apollo might be helpful. They concluded that triage could be valuable during high workload periods (when

multiple scans await review) but less relevant when only a few cases are pending, an early articulation that Apollo's value was situational and rhythm dependent, not a stable attribute of the model. In this setting, Apollo was immediately interpreted through existing prioritization practices. The neuroradiologist noted that triage colouring might be *"useful on busy shifts when you can't keep up"* but *"not on ordinary days with only one or two MR-brains waiting."* (Radiologist, Fieldnotes, 2022). This remark is analytically central because it shows how trust-in-use is tied to perceived necessity and relevance. The AI output is not evaluated as a freestanding truth claim; it becomes meaningful when it helps clinicians manage attention under pressure.

At the same time, the pilot revealed how easily help can become noise. The interface displayed urgent cases in red and occasionally continued flagging cases already reviewed, creating visual clutter. The neuroradiologist described the colour emphasis as stressful, prompting suggestions to replace colour overlays with arrows or simpler indicators. Here, the system's attempt to produce urgency could inadvertently produce uncertainty, precisely because urgency is already carefully managed through professional filtering and workflow rhythm. Across the pilot observations, another recurring issue was timing. When AI results appeared slightly too early or too late, they risked pulling attention at the wrong moment or arriving after the relevant decision had already been made. These micro-temporal misalignments were not merely usability bugs; they revealed how time mediates attention and accountability in imaging work. When timing aligned, the system could feel natural, almost invisible; when it did not, it demanded attention without offering actionable support. This is one place where early integration work becomes very concrete; infrastructuring is not only about technical connection or organizational approval, but also about whether a system can meet the temporal grain of the work it seeks to support.

By 2023, lessons from the early pilots were brought into a structured planning workshop for a new prospective study aimed at assessing whether Apollo could support triage and reduce length of stay for suspected stroke/TIA patients. Empirically, the meeting operated as both coordination and reflection between clinicians, innovation staff and developers. Participants revisited concrete frustrations from earlier pilots (e.g., alert interpretation, processing time, login/access friction) and used them to specify conditions under which the next study could be feasible and safe. Accountability concerns again took centre stage.

As one radiographer asked: *"If the neurologist acts on Apollo's flag and it's wrong, who carries that?"* (Radiographer, Prospective study planning workshop, Fieldnotes, 2023)

A senior neurologist responded: *“It can help us prioritize, but the decision stays with the doctor.”* (Neurologist, Prospective study planning workshop, Fieldnotes, 2023). Rather than handing off responsibility to the technology, this exchange shows accountability being actively reasserted and reconfigured around the system: Apollo could support sorting and attention, but responsibility remains located in clinical judgment. Interpretability was framed in the same way as a condition for professional accountability rather than a purely technical feature. Alongside these moral and professional negotiations, the workshop addressed infrastructural and organizational readiness.

The group made a deliberate scoping decision about what not to integrate yet. Full PACS integration was recognized as a long term goal, but it was not pursued for the trial to avoid complicating workflows. Instead, data flows and coordination were kept within a narrower clinical domain around the emergency neurology department. Finally, participants revisited performance-as-tempo (not performance-as-accuracy). A radiologist recalled earlier tests where processing time sometimes reached ten minutes; developers stated that the new setup consistently produced results within two minutes of scan completion. Plans were made for limited connectivity testing ahead of activation to verify stability under real conditions. The study design was discussed as a two-week alternating schedule (Apollo active one week and inactive the next) to create comparison periods and accommodate learning curves; data extraction plans included quantitative measures (e.g., length of stay, completion of stroke-related tasks) and qualitative observation of adaptation in practice.

## Discussion: Anticipatory infrastructuring and the work of making early clinical AI workable

This article contributes to the literature on the implementation of AI in healthcare by shifting attention from post-deployment adoption to the early, practical work through which AI output is made potentially workable within existing clinical infrastructures. First, it extends infrastructure scholarship on installed bases and relational integration (Star & Ruhleder, 1996; Star, 1999) by showing how feasibility work, pilots and pre-implementation planning are not merely preliminary stages, but sites where actors actively stabilise the conditions under which a system could later become part of routine work. Second, it extends research on infrastructuring as ongoing alignment work (Edwards et al., 2007; Pipek & Wulf, 2009) by demonstrating how early AI integration is shaped through negotiations of scope,

routability, tempo, visibility, and responsibility, long before a system becomes taken for granted in everyday practice.

We use anticipatory infrastructuring to interpret how actors work to stabilise the conditions under which an AI-based decision support system later could become part of routine clinical work. In the Apollo collaboration, this work shaped more than technical connectivity. It shaped possible future uses of the AI in this setting, what it could ask clinicians to notice, and how responsibility could remain intelligible when algorithmic suggestions appear during time-critical imaging work. The discussion develops the argument through the four moves identified in the findings. The point is not that these moves occur neatly one after another, but that they become tightly coupled. Small decisions about scope influence what counts as an acceptable interface. Decisions about visibility depend on tempo. Expectations for response depend on how reliable routing is, and on what the system is allowed to claim.

### *Delimiting scope and legitimacy: defining support that can be defended*

A first form of anticipatory infrastructuring concerns delimiting what the system may influence. We contribute to research on the material organisation of accountability in healthcare infrastructures (Bossen & Markussen, 2010) by showing how legitimacy is stabilised in advance through boundary making that constrains what an AI system is allowed to claim in order to prevent future responsibility regimes from becoming unworkable under acute conditions. In feasibility meetings and early pilot framing, participants repeatedly positioned the decision-support system as support for triage and attention, not as diagnostic authority. Statements such as *“an extra pair of eyes, but it shouldn’t speak for us”* and *“support decision-making, not decide”* do not only express preferences. They function as constraints that organise the collaboration. They shape what can be shown, what can be tested, and what can be justified to managers, legal staff, and clinicians.

This boundary work is anticipatory because it is oriented to the accountability pressures that routine use would bring. In acute imaging workflows, a tool that signals urgency easily becomes a demand for action. If alerts imply an obligation to verify or to document response, the tool can create additional work that is difficult to sustain under shift conditions. Participants’ reluctance toward constant alert verification and concerns about retrospective blame can be read as early attempts to prevent a future responsibility regime that clinicians regard as unrealistic. In this sense, legitimacy is stabilised by narrowing the system’s implied force. The system becomes *“acceptable”* by being framed as consultable input rather than an authority that generates new mandatory checks.

### *Stabilising routability and tempo: results have to arrive inside live decision windows*

A second form of anticipatory infrastructuring concerns routability and tempo. Here we contribute to infrastructure research on breakdown of anticipatory infrastructuring and installed bases (Star & Ruhleder, 1996; Star, 1999) by providing accounts of how routability and tempo become central conditions of actionability in early clinical AI integration, where late or unreliable outputs are not perceived as only technical failures but reshape whether AI can function as meaningful support. A clinical AI system's value in this case depends on whether outputs appear in time to matter for triage and acquisition-related decisions. The Apollo pilot breakdown around DWI transfer illustrates this clearly. Eleven cases were sent to the system, but results did not arrive. In that situation, the system's role cannot even be evaluated meaningfully, because the basic condition of actionability is not met. The practical conclusion, to rely on automated transfer until manual upload could be debugged, is therefore not a minor technical adjustment.

It is infrastructuring work that stabilises a plausible future use scenario, one where AI output can be treated as dependable enough to consult. The repeated attention to connectivity, expressed succinctly in the remark that *"it's never the scan itself that fails, it's always the connection,"* points to a general issue for early clinical AI. Performance claims are easy to make in environments where data are stable and time is not part of the evaluation. In stroke-ready imaging work, timing is part of what makes output meaningful. When results arrive too late, they move from "support" to "noise" or they become a retrospective signal that is risky rather than useful. Anticipatory infrastructuring in this case involved making tempo a central criterion of viability and treating the data pathway as part of the intervention rather than as background plumbing.

### *Tuning visibility in existing screen ecologies: avoiding a second queue*

A third form concerns visibility and attention. Here we contribute to research on infrastructuring and the relational visibility of infrastructure (Pipek & Wulf, 2009; Star, 1999) by showing how making AI "visible" is not a generic design problem, but a situated negotiation of how algorithmic output can be consulted within existing screen ecologies without producing a second queue of attention and coordination work. The pilot configuration that placed the clinical AI system Apollo on an additional screen is revealing because it shows that integration is not simply a matter of "embedding" AI into existing systems. It involves

negotiating how a new signal fits into established ways of maintaining overview. For neuroradiologists, workstation layouts support diagnostic concentration and sequencing of tasks. Shifting the system into browser tabs inside existing screens risks disrupting that organisation. Conversely, placing the system on a separate screen makes it consultable without displacing the established diagnostic ecology.

This is also where the system's interface design becomes consequential. Urgent colouring and the repeated flagging of already reviewed cases were experienced as stress-inducing and clutter-producing. In a workflow already saturated with signals and documentation pressures, a triage list can easily become a second queue that clinicians feel they have to monitor. The improvement suggestions around reducing visual stress and preventing already handled cases from reappearing indicate how trust-in-use is built in practice. It is not produced by asking clinicians to "trust" the tool. It is produced by making it possible to consult the tool without it demanding continuous attention, and without it undermining existing filtering practices that professionals use to protect focus.

### *Negotiating noticing and response: defining what non-action means*

A fourth form of anticipatory infrastructuring concerns responsibility expectations for noticing and responding. Here we contribute to research on pilots as socio-technical arrangements (Bansler & Havn, 2010) by demonstrating how expectations for noticing and responding to algorithmic suggestions are negotiated before routine use exists, shaping future accountability landscapes through mundane organisational and infrastructural decisions. The workshop question, "*If the neurologist acts on Apollo's flag and it's wrong, who carries that?*" makes visible a practical dilemma. A triage cue is only useful if someone is expected to see it and to treat it as meaningful. At the same time, in shift-based clinical work, expectations for response must remain compatible with staffing and role distribution. Anticipatory infrastructuring here consists in specifying what the AI output counts as in practice. Is it an optional cue that can be consulted when workload is high, or is it a signal that creates an obligation to respond regardless of context? That distinction shapes the future accountability landscape.

Importantly, this is not only about "who is responsible." It is about what kinds of work become necessary to make response expectations realistic. The planning decisions about onboarding, short instructional videos, and simplifying access via desktop shortcuts are part of this. They are mundane, but they determine whether noticing work can be sustained. If the system requires frequent logins, or if staff cannot reliably access it, the expectation to "respond" becomes empty.

In that sense, responsibility negotiations are inseparable from the infrastructural and organisational details that make response possible under real conditions.

The case addressed in this article helps clarify what is distinctive about anticipatory infrastructuring in early clinical AI integration. The work is not only about connecting a tool to a workflow. It is about making a future practice imaginable and defensible through concrete adjustments in scope, routing, tempo, visibility, and responsibility expectations. These adjustments are often decided before routine use exists, yet they shape what routine use can later become. This has implications for how early-stage assessment is framed. If assessments focus primarily on algorithmic accuracy, they risk missing the conditions that make output actionable. In time-critical imaging work, those conditions include whether the system can be kept within a legitimacy boundary that clinicians can defend, whether results arrive within live decision windows, whether visibility supports rather than fragments overview, and whether expectations for noticing and response are compatible with staffing and coordination realities. Treating these as central assessment concerns shifts attention away from whether “the model works” in the abstract, and toward whether the socio-technical arrangement can support meaningful use without producing new burdens or new accountability risks.

## Conclusion

This article has explored how early-stage integration of clinical AI can be understood as anticipatory infrastructuring, where clinicians, developers, and institutional actors negotiate how a system may attach to existing infrastructures, routines, and responsibility relations. Through a longitudinal rapid ethnographic study of feasibility work, pilots, and pre-implementation planning around Apollo in an acute stroke-ready MRI setting, the analysis has shown that making AI workable involves practical alignment work around scope, tempo, visibility, and responsibility. The concept of anticipatory infrastructuring offers a way of specifying how clinical AI becomes actionable before routine use is established. It highlights that integration is not only a matter of technical performance, but of shaping what a system is allowed to influence, when its outputs can matter, and how accountability remains intelligible within cooperative clinical practice. Early-stage integration thus emerges as a phase where future practice is gradually configured through organisational, infrastructural, and interpretive adjustments within an existing installed base. By foregrounding this alignment work, the

article contributes to discussions of how technologies become embedded in healthcare practice. It suggests that what later appears as “working AI” is stabilised through negotiations that precede adoption and that shape the conditions under which algorithmic output can function as meaningful support. Future research may examine how such anticipatory infrastructuring unfolds across other clinical settings and technologies, and how early alignment decisions influence longer-term routinisation and accountability arrangements.

## References

- Aanestad, M., Grisot, M., Hanseth, O., & Vassilakopoulou, P. (Eds.). (2017). *Information Infrastructures within European Health Care: Working with the Installed Base*. Springer. <https://doi.org/10.1007/978-3-319-51020-0>
- Bansler, J. P., & Havn, E. (2010). Pilot implementation of health information systems: Issues and challenges. *International Journal of Medical Informatics*, 79(9), 637–648. <https://doi.org/10.1016/j.ijmedinf.2010.05.004>
- Bossen, C., & Markussen, R. (2010). Infrastructuring and ordering devices in health care: Medication plans and practices on hospital wards. *Computer Supported Cooperative Work*, 19(6), 615–637. <https://doi.org/10.1007/s10606-010-9128-9>
- Edwards, P., Jackson, S., Bowker, G., & Knobel, C. (2007) “Understanding infrastructure: Dynamics, tensions, and design,” Report of a Workshop on “History & theory of infrastructure: Lessons for new scientific cyberinfrastructures.”
- Farič, N., Hinder, S., Williams, R., Ramaesh, R., Bernabeu, M. O., van Beek, E., & Cresswell, K. (2023). Early Experiences of Integrating an Artificial Intelligence-Based Diagnostic Decision Support System into Radiology Settings: A Qualitative Study. *Studies in Health Technology and Informatics*, 309, 240–241. <https://doi.org/10.3233/SHTI230787>
- Fox, S. E., Shorey, S., Kang, E. Y., Montiel Valle, D., & Rodriguez, E. (2023). Patchwork: The Hidden, Human Labor of AI Integration within Essential Work. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1), 1–20. <https://doi.org/10.1145/3579514>
- Helén, I., & Tarkkala, H. (2024). Adjusting expectations actionable: Personalised treatment plan in anticipation of data-driven healthcare. *Big Data & Society*, 11(4). <https://doi.org/10.1177/20539517241285384>
- Johansen, J.P., Almklov, P.G. & Mohammad, A.B. (2016). What can possibly go wrong? Anticipatory work in space operations. *Cognition, Technology & Work*, 18, 333–350. <https://doi.org/10.1007/s10111-015-0357-8>
- Karasti, H., Blomberg, J. (2018). Studying Infrastructuring Ethnographically. *Computer Supported Cooperative Work*, 27, 233–265. <https://doi.org/10.1007/s10606-017-9296-7>
- Khan, S. D., Hoodbhoy, Z., Raja, M. H. R., Kim, J. Y., Hogg, H. D. J., Manji, A. A. A., Gulamali, F., Hasan, A., Shaikh, A., Tajuddin, S., Khan, N. S., Patel, M. R., Balu, S., Samad, Z., & Sendak, M. P. (2024). Frameworks for procurement, integration, monitoring, and evaluation of artificial intelligence tools in clinical settings: A systematic review. *PLOS Digital Health*, 3(5), e0000514. <https://doi.org/10.1371/journal.pdig.0000514>
- Krag, C. H., Müller, F. C., Gandrup, K. L., Raaschou, H., Andersen, M. B., Brejnebo, M. W., Sagar, M. V., Bojsen, J. A., Rasmussen, B. S., Graumann, O., Nielsen, M., Kruuse, C.,

- & Boesen, M. (2023). Diagnostic test accuracy study of a commercially available deep learning algorithm for ischemic lesion detection on brain MRIs in suspected stroke patients from a non-comprehensive stroke center. *European Journal of Radiology*, 168, 111126. <https://doi.org/10.1016/j.ejrad.2023.111126>
- Lebovitz, S., Levina, N., & Lifshitz-Assaf, H. (2022). To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science*, 33(1), 126-148. <https://doi.org/10.1287/orsc.2021.1549>
- Malm-Nicolaisen, K., Ellingsen, G., Hertzum, M., Silsand, L., Severinsen, G.. (2024) Evolution of Information Infrastructures in Healthcare as Convergence of Digital Trajectories. *Computer Supported Cooperative Work*, 33, 1373–1400. <https://doi.org/10.1007/s10606-024-09502-0>
- McCraden, M. D., Joshi, S., Anderson, J. A., & London, A. J. (2023). A normative framework for artificial intelligence as a sociotechnical system in healthcare. *Patterns*, 4(11), 100864. <https://doi.org/10.1016/j.patter.2023.100864>
- Nielsen, C., Jensen, C. M., & Trettin, B. (2022). Epistemological aspects of participatory design studies conducted in the field of health science. *Journal of Advanced Nursing*, n/a(n/a). <https://doi.org/10.1111/jan.16293>
- Peek, N., Capurro, D., Rozova, V., & van der Veer, S. N. (2024). Bridging the gap: Challenges and strategies for the implementation of artificial intelligence-based clinical decision support systems in clinical practice. *Yearbook of Medical Informatics*, 33(1), 103–114. <https://doi.org/10.1055/s-0044-1800729>
- Pipek, V., & Wulf, V. (2009). Infrastructuring: Toward an integrated perspective on the design and use of information technology. *Journal of the Association for Information Systems*, 10(5). <https://doi.org/10.17705/1jais.00195>
- Ramsay, A. I. G., Crellin, N., Lawrence, R., Walton, H., Bagri, S., Dodsworth, E., Elphinstone, H., Gleeson, F., Halliday, A., Herbert, K., Lloyd, J., Massou, E., Mehta, R., Morris, S., Ng, P. L., O'Regan, T., Sherlaw-Johnson, C., & Fulop, N. J. (2025). Procurement and early deployment of artificial intelligence tools for chest diagnostics in NHS services in England: A rapid, mixed method evaluation. *eClinicalMedicine*, 89. <https://doi.org/10.1016/j.eclinm.2025.103481>
- Star, S. (1999). The ethnography of infrastructure. *American Behavioral Scientist*, 43(3), 377–391. <https://doi.org/10.1177/00027649921955326>
- Star, S., & Ruhleder, K. (1996). Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. *Information Systems Research*, 7(1), 111–134. <https://doi.org/10.1287/isre.7.1.111>
- Steinhardt, S. B., & Jackson, S. J. (2015). Anticipation Work: Cultivating Vision in Collective Practice. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, 443–453. <https://doi.org/10.1145/2675133.2675298>
- Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Geerts, B., Ibrahim, M., Liu, X., Mateen, B. A., Mathur, P., McCraden, M. D., Morgan, L., Ordish, J., Rogers, C., Saria, S., Ting, D. S. W., McCulloch, P. (2022). Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *BMJ*, 377, e070904. <https://doi.org/10.1136/bmj-2022-070904>
- Vindrola-Padros, C., & Vindrola-Padros, B. (2018). Quick and dirty? A systematic review of the use of rapid ethnographies in healthcare organisation and delivery. *BMJ Quality & Safety*, 27(4), 321–330. <https://doi.org/10.1136/bmjqs-2017-007226>

# AI som beslutningsstøtte til fraktur i skadestuen?

## Billeddiagnostik, tidslighed og patientforløb i afprøvningen af fraktur-AI

Michal Frumer<sup>1,2</sup>

Christina Egelund Antonsen<sup>3</sup>

Maja Hojer Bruun<sup>4</sup>

---

<sup>1</sup>Forskningsenheden, Regionshospitalet Horsens

<sup>2</sup>Interacting Minds Centre, Institut for Klinisk Medicin, Aarhus Universitet

<sup>3</sup>AUH Innovation & International Affairs, Aarhus Universitetshospital

<sup>4</sup>DPU – Danmarks Institut for Pædagogik og Uddannelse, Aarhus Universitet

*micfru@rm.dk*

*christina.egelund.antonsen@rm.dk*

*mhbbruun@edu.au.dk*

Frumer, Michal; Antonsen, Christina Egelund & Bruun, Maja Hojer. 2026. AI som beslutningsstøtte til fraktur i skadestuen? Billeddiagnostik, tidslighed og patientforløb i afprøvningen af fraktur-AI. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 48-75. DOI: 10.7146/TFSS.V25I44.156610

*Indsendt 04/25, accepteret 01/26, udgivet 06/26*

*Denne artikel undersøger afprøvningen af artificial intelligence (AI) i to danske skadestuer, hvor såkaldt fraktur-AI markerer mulige knoglebrud og ledskred på røntgenbilleder af blandt andet arme og ben. Baseret på etnografisk feltarbejde og interviews analyserer vi, hvordan AI som beslutningsstøtte praktiseres, fortolkes og anvendes i den kliniske hverdag. Vi argumenterer for, at fraktur-AI ikke blot er et teknisk redskab, men en del af*

*et socioteknisk system, hvor teknologiens betydning og anvendelse formes af skadestuens tidlige logikker, organisatoriske rytmer og politiske forventninger. Analysen identificerer tre centrale tidsligheder i samspillet mellem AI og klinisk praksis: punktuering af tid, hvor diagnostiske processer forenkles til en binær vurdering af brud/ikke-brud; tempo, hvor AI legitimeres gennem forestillinger om effektivitet og optimering af ventetid; og AI-som-løfte, hvor forventninger til fremtiden bliver betydningsbærende trods manglende nutidig evidens.*

*Empirisk bidrager artiklen med indsigt i, hvordan AI-teknologier i praksis påvirker diagnostiske arbejdsgange og formes af og indgår i eksisterende faglige, organisatoriske og politiske interesser. Teoretisk argumenterer vi for, at AI som beslutningsstøtte adskiller sig fra tidligere algoritmiske styringsredskaber ved at mangle forankring i proces og klinisk kontekst, hvilket udfordrer selve forståelsen af beslutningsstøtte. Hermed tilbyder vi et analytisk blik på AI som en teknologi, der gør fremtiden styrende for nutidens praksisser og skaber nye former for organisatorisk og professionel autoritet i sundhedsvæsenet.*

## AI as decision support for fractures in the Emergency Department? Diagnostic imaging, temporalities, and patient pathways in the testing of fracture-AI

*In this article, we explore the trial implementation of artificial intelligence (AI) in two Danish emergency departments, specifically a system called “fracture-AI” that marks potential fractures and dislocations on X-ray images of arms, legs, and other areas. Based on ethnographic fieldwork and interviews, we analyse how AI, designed as a decision-support tool, is practiced, interpreted, and applied in everyday clinical work. We argue that fracture-AI is not merely a technical instrument but part of a socio-technical system in which the technology’s meaning and use is shaped by the temporal logics of the emergency department, organizational rhythms, and political expectations. The analysis identifies three central temporalities in the interplay between AI and clinical practice: punctuation of time, where diagnostic processes are reduced to a binary assessment of fracture/no-fracture; tempo, where AI is legitimized through imaginaries of efficiency and reduced waiting time; and AI-as-promise, where expectations of future benefits become performative despite a lack of present clinical evidence.*

*Empirically, the article contributes insights into how AI technologies in practice affect diagnostic practices and are shaped by and embedded in existing professional, organizational, and political interests. Theoretically, we argue that AI as decision support differs from earlier algorithmic support tools by lacking anchoring in process and clinical context, thereby challenging the very notion of “decision support”. In doing so, we offer an*

*analytical perspective on AI as a technology that renders the future operative in the present and generates new forms of organizational and professional authority within healthcare.*

## Introduktion

Den unge ortopædkirurgiske læge, Fie, som sidder ved det lille bord overfor én af denne artikels forfattere, har en blå hospitalsuniform på. Blusen er en klar blå, løstsiddende og kortærmet busseronne med V-hals til at trække over hovedet. Det er praktisk og hygiejnisk arbejdstøj, som bruges på skadestuen ligesåvel som på operationsstuerne. Lægen har stadig sin lyseblå operationshue på, som dækker hendes opsatte hår. Hun kommer direkte fra en ortopædkirurgisk operation, men har lige nået at spise lidt mad inden vores interview. Fie er ansat på et af de danske hospitalet, som i perioden 1. september 2024 til 28. februar 2025 har afprøvet to konkrete artificial intelligence (AI) baserede teknologier, der kan markere knoglebrud og ledskred på røntgenbilleder af bevægeapparatet (f.eks. hånd, fod og albue).

Få minutter inde i samtalen, efter at Fie har fortalt om sit speciale i ortopædkirurgi, spørger interviewereren, om Fie vil fortælle lidt om sine erfaringer med det nye AI-redskab. Fie fortæller, at det ikke umiddelbart har ændret hendes arbejdsgang med røntgenbilleder, og hun er heller ikke overbevist om, at det er en hjælp. Indimellem er det mere en distraktion. Fie opsummerer:


*I 80% af tilfældene så forholder jeg mig næsten ikke til den. Altså jeg forholder mig til, at den er der. Og jeg konstaterer, at AI'en har sagt noget. Men så går jeg tilbage til min kliniske beslutningstagning.*

I denne artikel undersøger vi afprøvningen af fraktur-AI<sup>1</sup> til beslutningsstøtte i skadestuen i forbindelse med simple knoglebrud og ledskred i arme og ben. Nye AI-redskaber bliver i disse år afprøvet i de fem danske regioner, der styrer det danske sundhedsvæsen, og nogle steder også implementeret som "digital faglig beslutningsstøtte" (Dansk Selskab for Patientsikkerhed, 2023). Gennem årene er mange forskellige beslutningsstøtteteknikker, såsom kliniske protokoller, blevet indført som del af opbygningen af evidensbaseret medicinsk praksis (Berg, 1997). Beslutningsstøtte med fraktur-AI betyder i denne afprøvning, at AI-redskabet

---

1. Vi anvender samlebetegnelsen "fraktur-AI" om de to specifikke AI-redskaber, der blev afprøvet i skadestuerne. Begge redskaber, som var udvalgt af regionens enhed for Indkøb og Medicoteknik og blev præsenteret af produkternes danske leverandører, baserer sig teknisk på deep learning og neurale netværk. Under feltarbejdet og interviews blev de afprøvede teknologier omtalt som "frakturalgoritmen" og "AI'en".

analyserer røntgenbilleder og markerer med en firkant, hvor der sandsynligvis er brud eller ledskred (Figur 1). Læger som Fie kan herefter vurdere disse informationer i lyset af deres egne observationer og tolkning af røntgenbilledet og tage stilling til, om der er tale om et brud eller ej. Det er uanset hvad altid en læge, som beslutter, hvad næste skridt skal være: yderligere undersøgelser af patienten, behandling eller afslutning af patientens forløb.



*Figur 1: AI-firkanter på røntgenbilleder. Eksempel fra én af leverandørerne*

I artiklen forstår vi AI og fraktur-AI som mere end et teknisk redskab eller bestemt program, nemlig som et socioteknisk system af menneske-teknologi-relationer, der også omfatter alle de forskellige former for viden, organisering og kultur, som får systemet til at fungere (Pfaffenberger, 1992). At betragte teknologier som sociotekniske systemer indebærer, at de ikke fungerer ens på tværs af kontekster. Det samme redskab eller den samme algoritme kan udfolde sig – eller *enactes*, som man ville sige på engelsk – på forskellige måder i forskellige kulturelle og organisatoriske sammenhænge. Når de anvendes af forskellige praktikere, bliver teknologierne dermed faktisk til noget forskelligt hvert sted (Seaver, 2017). Inden for medicinsk praksis er forskellige protokoller og computerbaserede systemer til beslutningsstøtte blevet indført i løbet af det 20. århundrede og op til i dag, i takt med at lægekunst blev transformeret til moderne videnskab og evidensbaseret praksis (Berg, 1997). I skadestuerne ser vi mange forskellige beslutningsstøtteredskaber og teknikker, f.eks. til diagnosticering eller risikostratificering, der strukturerer medicinsk praksis på bestemte måder. I den sammenhæng nøjes fraktur-AI og andre former for AI ikke med at sortere informationer og strukturere forskellige arbejdsgange, men sigter også mod at automatisere eller erstatte arbejdsgange. Vores perspektiv på tekniske redskaber som værende en del af sociotekniske

systemer muliggør en undersøgelse af de forskellige vidensformer, løfter og tidsligheder, som knytter sig til anvendelsen af fraktur-AI som et beslutningsstøtteredskab. Det er med andre ord ikke blot egenskaber ved fraktur-AI som teknisk redskab, som muliggør, at denne teknologi har fundet sin vej ind i skadestuen og kan fungere som beslutningsstøtte dér, men også løfter om fremtidige forbedringer af AI, forestillinger om optimering af arbejdsgange og den plads, som fraktur-AI får tildelt i koordineringen af arbejdet i skadestuen.

Når vi skal forstå den praksis og de løfter, som knytter sig til billeddiagnostik med AI, må vi udvide vores blik ikke kun til radiologers interaktioner med AI, men også brugen af fraktur-AI i andre hospitalsafdelinger, og til der hvor teknologiudviklere og -leverandører, politikere og ledere forestiller sig, at AI kan effektivisere ved at aflaste eller ligefrem erstatte radiologer (Topol, 2019; Watson & Wozniak-O'Connor, 2025). Med udgangspunkt i fraktur-AI som et socioteknisk system, flytter vi fokus fra radiologers arbejde alene til et bredere tværfagligt og organisatorisk perspektiv, hvor vi undersøger, hvordan AI indgår i samarbejdet mellem læger på tværs af specialer. Fraktur-AI er her et eksempel på en fremtrædende, nyere form for AI, der anskues som digital faglig beslutningsstøtte inden for billeddiagnostik. Vores case gør det yderligere muligt at undersøge AI i en kompleks og tidsfokuseret organisation med en høj grad af specialisering, arbejdsdeling og tværfagligt samarbejde. Vi beskriver disse forhold omkring vores studie nærmere i artiklens første tre afsnit om hhv. AI i billeddiagnostik blandt radiologer, metode og skadestuen som felt.

Med udgangspunkt i Fie og andre lægers måde kort at forholde sig til AI-redskabets markeringer, før de vender tilbage til det, de kalder deres "kliniske beslutningstagning", argumenterer vi i denne artikel for, at fraktur-AI ikke blot påvirker diagnostiske arbejdsgange, men også grundlæggende former skadestuens tidsligheder. "Klinisk beslutningstagning" er en tidsligt udstrakt proces (Frumer, 2023), og i artiklen udforsker vi forholdet mellem AI og tidslighed ud fra tre temaer: Først viser vi, hvordan teknologien medvirker til en fastfrysning og 'punktuering af tid', idet fraktur-AI forstærker røntgenbilledet som et isoleret datapunkt og øjebliksbillede, som lægger op til en binær vurdering af brud/ikke-brud og overskygger komplekse tidslige vekselvirkninger mellem fortid, nutid og fremtid i den diagnostiske proces.

Derefter fokuserer vi på 'tempo og ventetid', hvor teknologien indskrives i forestillinger om accelererede beslutningsgange og effektivisering, men hvor idealet om tidsbesparelse modsiges af de faktiske rytmer i klinisk praksis.

Til slut diskuterer vi 'AI-som-løfte', hvor afprøvningen af fraktur-AI viser sig båret af et særligt fremtidsblik: Selvom tilliden til redskabet i nutiden er begrænset,

opretholdes projektet og fraktur-AI's legitimitet gennem performative løfter om en anden, endnu-ikke-realiseret tidslighed. Disse forventninger opretholdes under afprøvningen og deles på forskellig vis af projektdeltagere, leverandører og klinikere, der ser potentialet for fremtidig nytte. Artiklen bidrager dermed til forståelsen af fraktur-AI som et socioteknisk system, der former klinisk praksis og tidsligheder samt udfordrer eksisterende forestillinger om beslutningsstøtte.

## AI i billeddiagnostik: Fokus på radiologers arbejde

AI-teknologier anses af politikere og leverandører for at have stort potentiale til at effektivisere sundhedsvæsenet ved f.eks. at muliggøre hurtigere og mere præcise vurderinger. Særligt inden for billeddiagnostik er AI blevet udviklet, afprøvet og implementeret, og i USA blev cirka 76% af alle AI-teknologi-godkendelser (777 ud af 1016) i 2024 givet til radiologiske teknologier (US Food & Drug Administration, 2025). Også i Danmark fremhæves billeddiagnostiske AI-teknologier som konkrete AI eksempler inden for sundhedsområdet f.eks. til at vurdere røntgenbilleder, hvor et konkret redskab forventes at kunne reducere både ventetid for patienten og mængden af klagesager (Regeringen, 2024, s. 7). I februar 2025 afsatte Regeringen, Danske Regioner og KL som led i en større AI-satsning på sundheds- og ældreområdet betydelige midler til fraktur-AI, der dermed bliver en hjørnesten i den danske AI-strategi de kommende år (Danske Regioner, 2025).

Der er flere grunde til det tætte parløb mellem AI og radiologi. Radiologi har historisk været et teknologisk medieret område, hvor læger løbende har indarbejdet nye teknologier og arbejdsformer (Pesapane et al., 2018; Saunders, 2008). Siden 1970'erne er billeddiagnostik blevet digitaliseret, f.eks. med introduktion af CT-skanninger i 1972 (McCollough & Rajiah, 2023), og fra årtusindskiftet blev almindelige røntgenundersøgelser og arbejds gange digitaliseret (Saunders, 2008). Den digitale dataproduktion har ikke alene muliggjort træning af AI-modeller, men også fremmet udviklingen af mønstergenkendelse og machine learning til identifikation af knoglebrud, vævsændringer og andre billedmæssige træk, der tolkes som relevante. Derudover drives brugen af AI i radiologien af en massiv stigning i brugen af billeddiagnostik og en betydelig mangel på radiologer (Kidwai, 2023; Olsen et al., 2023).

Den medicotekniske fortælling fremhæver, at AI er i gang med hurtigt at transformere radiologien eller ligefrem at revolutionere diagnostisk præcision og klinisk effektivitet i billeddiagnostikken (Katal et al., 2024; Khalifa & Albadaawy, 2024). Dette indebærer en ændring i forståelsen af radiologers professionelle

ekspertise fra at være centreret om lægers "subjektive, perceptuelle færdigheder" til at blive "en objektiv videnskab" (Pesapane et al., 2018). AI antages at have potentiale til at overtage mange rutinemæssige detektions-, karakteriserings- og kvantificeringsopgaver, dog uden at erstatte radiologer helt. Flere og flere videnskabelige artikler er optagede af særligt radiologers forventninger og holdninger til AI og forsøger at komme med retningslinjer for udvikling og implementering af AI (Huisman et al., 2021; Katal et al., 2024).

Den kvalitative litteratur om billeddiagnostisk AI har ligeledes primært fokus på radiologers praksis. Studier viser, hvordan AI både skaber nye former for klinisk blik og nye arbejdsopgaver (Anichini et al., 2024; Lebovitz et al., 2022; Lombi & Rossero, 2024). Lebovitz et al. (2022) beskriver f.eks., hvordan radiologer enten ignorerede, ukritisk accepterede eller gennem "AI-afhøringspraksisser" aktivt forhandlede med AI-resultaterne. I dansk sammenhæng fremhæves radiologens rolle som "kvalitetskontrollør" og som den, der afgør, om resultaterne overhovedet er meningsfulde (Kristensen, 2025). Samarbejdet med AI indebærer altså både forhandlinger om ekspertise og forskydninger af arbejdsopgaver.

I denne artikel vender vi blikket mod AI som beslutningsstøtte dér, hvor radiologerne ikke er, i dette tilfælde i skadestuen. Vi undersøger, hvordan fraktur-AI ikke blot påvirker kliniske beslutninger eller opgaver, men også former skadestuens tidsligheder. Hvor tidligere studier har fokuseret på radiologer, der sidder afsondret fra patienter og arbejder med billeder som abstrakte data, bidrager vores undersøgelse med indsigt i, hvordan AI indlejres i en klinisk praksis, hvor skadestuelæger har direkte kontakt med patienter og deres kroppe, må forholde sig til patienter i venteværelser og navigere komplekse patientflow og tempi.

## Metode og empirisk materiale

Vores analyse bygger på kvalitativ, etnografisk empiri genereret på to danske hospitaler i løbet af afprøvningen af fraktur-AI den 1. september 2024 til 28. februar 2025 og ved deltagelse i projektmøder op til, under og efter afprøvningen. Afprøvningen foregik på skadestuerne på Aarhus Universitetshospital (AUH) og Regionshospitalet Horsens (RHH), men var foranlediget af ledelsen i Region Midt i samarbejde med regionens enhed for Indkøb og Medicoteknik. Forfatternes adgang til felten blev skabt gennem afprøvningsprojektets aftale med skadestuerne om at gennemføre en undersøgelse af afprøvningen. Selvom projektet internt gik under navnet "afprøvning af fraktur-algoritme", omfattede det reelt to forskellige teknologier fra separate leverandører: RBfracture fra Radiobotics på den ene

skadestue og Techcare Alert fra Milvue (formidlet af Human Bytes) på den anden<sup>2</sup>. Vi har valgt at benævne de to teknologier samlet som "fraktur-AI", da vores analyse ikke har til formål at sammenligne redskaberne. De to systemer blev i praksis integreret i skadestuernes arbejdsgange på tilsvarende vis, og de relaterede kliniske og organisatoriske processer var sammenlignelige, hvilket gør én samlet analyse metodisk meningsfuld. Denne samlede betegnelse beskytter samtidig deltagerne i vores undersøgelse, som derved ikke indirekte kan identificeres i forhold til hvilket hospital, de er ansat på.

Frumer, Antonsen og Bruun samarbejdede om at udføre semistrukturerede interviews, deltagerobservation og uformelle samtaler på de to skadestuer og til projektmøder for projektet "afprøvning af fraktur-algoritme". Med en fælles interviewguide gennemførte vi 32 semistrukturerede interviews med klinikere på tværs af specialer og erfaringsniveau. Grundet de to hospitalers organisatoriske forskelle, repræsenterer interviewpersonerne på RHH forskellige afdelinger (se tabel 1), mens der på AUH primært har været fokus på skadestuens personale (herunder ortopædkirurger) men med længere etnografiske samtaler med de øvrige specialafdelinger.

**Tabel 1:** *Oversigt over interviews*

	Skadestue (akutafdeling)	Ortopædkirurgi	Radiologi
<b>Regionshospitalet Horsens 14 interviews</b>	7 yngre læger	2 yngre læger 2 erfarne læger	1 yngre læge 2 erfarne læger
<b>Aarhus Universitetshospital 18 interviews</b>	10 yngre læger 7 erfarne læger (herunder ortopædkirurger) 1 behandler- sygeplejerske	Samtale med 1 erfaren læge	Deltagerobservation og samtale med 1 erfaren læge

2. Der findes ikke en videnskabelig publikation som på sammenlignelige, danske "real-life" data vurderer de to algoritmers performance, og vi har derfor ikke inkluderet en sådan teknisk beskrivelse, men henviser til leverandørerne, hvis yderligere information har interesse.

Alle interviews blev foretaget på hospitalerne på dansk og blev optaget og transskriberet. På AUH blev otte interviews foretaget af en studentermedhjælper. I interviewene fik deltagerne mulighed for at beskrive deres arbejds erfaring og arbejdsopgaver, hvorefter interviewene kredsede om konkret brug af fraktur-AI med dybdegående beskrivelser af arbejdspraksis og specifikke cases.

Foruden interview og deltagelse i projektmøder har Bruun<sup>3</sup> haft seks dage med deltagerobservation i november 2024 til januar 2025 på AUH med udgangspunkt i skadestuens kliniske dagligdag (fire dagvagter og to nattevagter). Disse dage bestod primært af observationer af lægers og sygeplejerskers arbejde i skadestuen og interaktion med fraktur-AI, herunder uformelle samtaler om diagnostik og brug af AI. Som del af deltagerobservationen fulgte Bruun patientforløb og røntgenbilleders vej gennem hospitalet fra røntgenafdelingen til skadestuen og røntgenkonferencen til traumekonferencen i radiologisk afdeling.

De konkrete interviews og deltagerobservation blev aftalt lokalt og med respekt for deltagerens tid, interesse og mulighed for at sige nej. Interviewpersoner og personale blev mundtligt informeret om det, som gik under betegnelsen "den kvalitative undersøgelse af fraktur-AI-afprøvningen". Det empiriske materiale er indsamlet og opbevaret med henvisning til GDPR-forordningens forskningshjemmel og behandles anonymiseret og med fortrolighed, da der bl.a. er tale om interview og samtaler med yngre læger under uddannelse og under andres ledelse i hierarkiske organisationer. Navne er pseudonymiseret, og citater er redigeret for sproglig forståelse. Der er ikke systematisk indsamlet materiale om patienter.

Sammen analyserede vi med fokus på indhold og mønstre på tværs af materialet i en kollektiv, iterativ proces med i alt fire workshops og en lang række møder. Vi identificerede først en række foreløbige temaer og sensibiliserende begreber i materialet, som vi efterfølgende nærlæste materialet med og som orienterede vores videre analyse og dataindsamling (Hammersley & Atkinson, 2019). Herefter kodede vi et udvalgt materiale, diskuterede yderligere kodning og lagde os fast på temaet og argumenterne om tidslighed til denne artikel. Selvom kun et udvalg af materialet blev systematisk kodet, blev hele materialet inddraget gennem fælles gennemlæsninger, analytiske workshops og iterative diskussioner, hvor vi afprøvede og justerede vores forståelser af mønstre, afvigelser og tematiske sammenhænge. Skriveprocessen med konkrete ideer, udkast og kommentering var en del af den analytiske proces, hvor de empiriske eksempler og analytiske pointer om tidslighed løbende blev afgrænset og tilspidset.

---

3. Bruuns forskning er støttet af DFF, grant ID 10.46540/2063-00039B

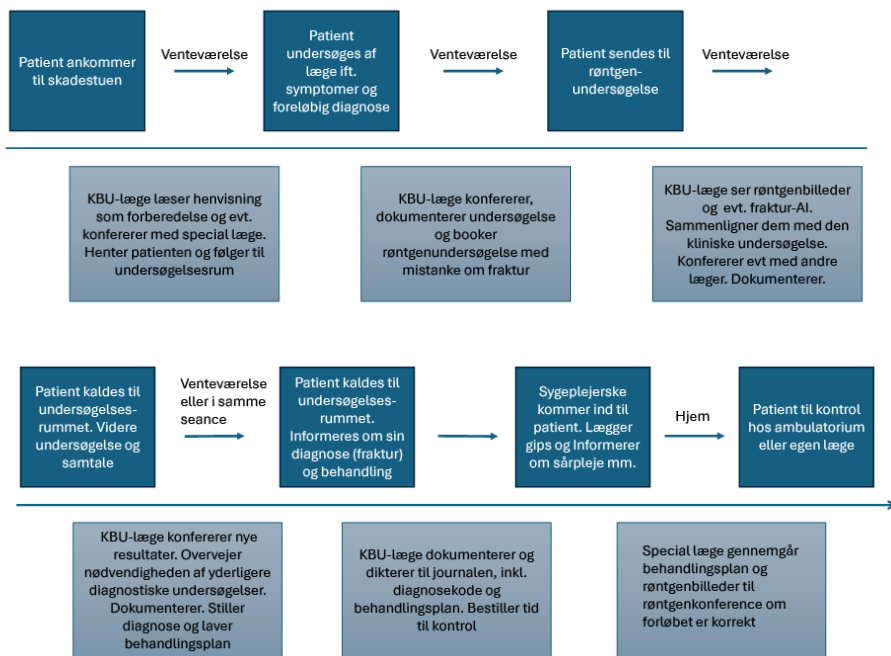
## Skadestuen som felt

Når det gælder den konkrete organisering af de danske skadestuer, herunder røntgen og skanning, er der store forskelle. Et eksempel er bemanningen af skadestuerne og den organisatoriske forankring af personalet. På AUH og RHH er begge skadestuer forankret i hospitalernes akutafdelinger, dog med nogle væsentlige forskelle. AUH er Danmarks største hospital med landets største skadestue. Her er ortopædkirurgiske speciallæger ansat i akutafdelingen med vagter i skadestuen. Ortopædkirurger er specialister i sygdomme i bevægeapparatet, herunder sygdomme i knogler, led og muskler. RHH er til gengæld et af regionens mindste akuthospitaler og "låner" ortopædkirurger fra hospitalets ortopædkirurgiske afdeling. Det betyder konkret, at de superviserende speciallæger i skadestuen er ansat i ortopædkirurgisk afdeling, men at akutafdelingen betaler for deres løn efter en fastsat norm og samarbejdsaftale. Først i 2017 blev akutmedicin godkendt som et selvstændigt lægefagligt speciale, og en ny akutlægeuddannelse blev oprettet i 2019. De første akutlæger er ved at blive færdiguddannet, og med tiden skal læger med dette speciale ansættes i landets akutafdelinger og skadestuer. Nogle hospitaler oplever i højere grad en mangel på læger til at bemane skadestuer og skadesklinikken. I nogle af landets skadestuer kan "yngste" vagthavende læge, som tager imod og undersøger patienterne først, være medicinstuderende, og de erfarne læger kan være optagede af andre opgaver såsom operationer eller være på telefonisk kald. Der kan endvidere gå lang tid, inden en radiolog har beskrevet røntgenbillederne. Af den grund kan de yngre læger ofte stå alene med beslutninger om diagnose og behandling af den form for brud, som fraktur-AI er udviklet til.

Diagnosticering og behandling af patienter kræver forskellige former for samarbejde og konferering, som typisk foregår gennem personlig samtale, opringninger eller udveksling af informationer gennem den elektroniske patientjournal (EPJ), se skitseringen i figur 2.

Når patienten har været til røntgen, kan røntgenbillederne åbnes af skadestuelægerne i EPJ. Beskrivelser og analyser af røntgenbillederne med billeddiagnostiske fund kommer lidt senere fra radiologer eller beskrivende radiografer i radiologisk afdeling. I praksis betyder dette, at lægerne i skadestuen ofte gennemgår røntgenbillederne og sammenholder dem med deres fund fra den fysiske undersøgelse af patienten, før de modtager de radiologiske fund. I nattetid prioriterer radiologer og beskrivende radiografer typisk CT-, MR- og ultralydbilleder over almindelige røntgenbilleder fra skadestuerne, men beskrivelserne foreligger senest næste hverdag, hvor de kan ses i EPJ. Nogle

steder har manglen på radiologer ført til, at radiologiske beskrivelser foretages af telemedicinske klinikker af radiologer i andre lande (f.eks. Albinus, 2022). Alle diagnostiske og behandlingsmæssige beslutninger, som er foretaget i skadestuen, gennemgås og sammenholdes med røntgenbillederne af en speciallæge først-kommende hverdag ved en røntgenkonference, hvor eventuelle fejl kan påpeges og rettes. I nærværende artikel er det primært skadestuelægernes vurdering af røntgenbilleder, som vi beskæftiger os med. Det er uvist, om det er til dem eller til radiologerne, at fraktur-AI er udviklet, men det er ikke-radiologers brug af fraktur-AI, som vi fokuserer på her.



Figur 2: Eksempel på en diagnose- og behandlingsproces i skadestuen

## Punktuering af tid i den diagnostiske proces

For at analysere hvordan fraktur-AI praktiseres som beslutningsstøtte i skadestuen, beskriver vi først den diagnostiske proces. Med dette som udgangspunkt diskuterer vi, hvordan tidsligheder i diagnose og behandling formes, når fraktur-AI anvendes som beslutningsstøtte. Vi argumenterer for, at der sker en

punktuering af tid i datapunkter, hvilket giver røntgenbilledet øget vægt og skygger for de vekselvirkende diagnostiske og behandlingsmæssige tidsligheder i den kliniske praksis.

*Skadestuens vagtrum summer af samtaler på kryds og tværs af rummet en helt almindelig tirsdag formiddag i januar 2025. Rasmus, som har været KBU-læge<sup>4</sup> siden 1. august, ser på sin computer sammen med antropologen. Han har lige undersøgt en ung mand, som henvendte sig til skadestuen med ondt i hånden efter et fald. Rasmus fandt såkaldt trepunktsømhed, som er tegn på brud af bådbenet i håndroden (scaphoideum på latin), og nu er billederne kommet fra røntgen. Scaphoideum er en meget hyppig fraktur, som dog kan være svær at se på røntgenbilleder, forklarer Rasmus. Med en hurtig bevægelse, som vi tit så blandt de yngre læger, når de åbnede røntgenbilleder med AI-analyse, skærmer Rasmus med hånden for nederste højre hjørne af sin computerskærm. Han skynder sig at klikke videre til de enkelte projektioner, som viser hånden oppefra, fra siden og skråt men uden AI-markeringer. "Jeg prøver at dække lidt til med hånden for at få læringen med, så jeg ikke ser resultatet med det samme. Læringsmæssigt kan den [fraktur-AI] godt snyde én og gøre én lidt doven. De havde den ikke, da jeg startede [på skadestuen], og der kiggede jeg måske lidt grundigere på billederne."*

*Rasmus kalder på Morten, som taler med en kollega i den anden ende af rummet. Han hvisker til antropologen: "Morten er en garvet ortopædkirurg." "Morten, har du tid til at kigge på en scaphoideum?" Morten kan ligesom Rasmus ikke se nogen fraktur på røntgenbilledet. Han spørger Rasmus: "Er der klinik nok?", hvormed han mener, om Rasmus' kliniske undersøgelse af patienten viser tegn på en fraktur. Morten: "AI'en markerer nogle gange de lidt bløde skygger her som fraktur." Men AI'en har ikke markeret nogen fraktur, og de bliver enige om, at der kun er tale om en forstuvning.*

*Da Morten er gået, åbner Rasmus en ny projektion, som han kalder scaphoideum-projektionen, og som han helt havde glemt, siger han til antropologen. Efter et øjeblik ringer han til Morten: "Nu fik jeg slet ikke vist dig scaphoideum-billedet, og her siger AI, at det ER en fraktur." Morten: "Er der klinik til det?" I løbet af en kort samtale bliver Rasmus og Morten enige om at behandle mandens hånd, som om der er en fraktur på scaphoideum for en sikkerheds skyld. Rasmus forklarer: "Jeg har jo ikke ret til selvstændigt virke, men nu er den hånd da konfereret. Nu behandler vi det som et*

---

4. En KBU-læge er en nyuddannet læge under klinisk basisuddannelse, et obligatorisk praktikforløb, hvor lægen endnu ikke må arbejde fuldstændigt selvstændigt og derfor træffer beslutninger under vejledning af mere erfarne kollegaer.

*brud, og han får han gips på. Om to uger kommer han så ind igen, og hvis der stadig er trepunktsømhed, så laver vi en MR-scanning. Det er det nye [behandlings]regime. Ellers kunne vi hele tiden lave MR-scanninger."*

*Rasmus vender sig om til en af sygeplejerskerne og forhører sig, hvilken gips hun vil anbefale og beder hende om at lægge den.*

Som denne episode viser, er den diagnostiske proces udstrakt over tid og rekursiv forstået på den måde, at læger vender tilbage til deres patienter og røntgenbilleder flere gange og konfererer med kollegaer for at stille diagnoser, som ændres og tilpasses henad vejen. Lægerne danner sig ved de første undersøgelser et billede af, hvad patienterne kan fejle, og i dette tilfælde førte det til bestilling af røntgenbilleder med nogle bestemte projektioner. Hvis "klinikken passer", og røntgenfund stemmer overens med den kliniske undersøgelse, kan en behandlingsplan lægges med det samme. Andre gange må lægen tilbage til patienten og foretage yderligere undersøgelser, finde ældre røntgenbilleder frem fra tidligere skadestuebesøg eller bestille andre diagnostiske undersøgelser. Den diagnostiske proces er med andre ord præget af bevægelser, som både orienterer sig bagud i tid (hvad ledte op til skaden), mod nuet (hvordan har patienten det nu) og fremad (forventning om bestemte skader og behandlinger). Disse tidlige orienteringer er ikke klart adskilte, men kendetegnet ved vekselvirkning mellem de forskellige tidlige orienteringer f.eks. i relation til en vurdering af smerte.

Røntgenbilleder kan ligesom journalnotater, EPJ og lignende teknologier betragtes som en informationsinfrastruktur med klassifikations- og standardiserings-systemer, som får den diagnostiske proces til at glide ved skabe en tidslig orden og samtidig antager en mere tidløs, dekonstekstuel form (Star, 1999; Bowker & Star, 1999). Røntgenbilleder forstået som infrastruktur i det diagnostiske arbejde har en tidslig rækkevidde, som rækker ud over klinikernes praksis på stedet (cf. Star, 1999, s. 381). Man kan sige, at den kliniske praksis omkring diagnosticering og behandling, som ikke foregår lineært men rekursivt, og som er kontekst- og situationsbunden (Suchman, 1987), bliver sat ind i en lineær, journaliserbar og overskuelig orden gennem fastfrysninger af tid i informationsinfrastrukturer som røntgenbilleder. I klinikken fletter den situerede praksis-tid og røntgenbilleders fastfrosne tid sig ind i hinanden.

I vores materiale blev det imidlertid tydeligt, at fraktur-AI skabte særlige ophold og forstyrrelser i den praktiske kliniske proces omhandlende diagnostik og behandling, hvor lægernes opmærksomhed blev trukket væk fra de kliniske vurderinger og yderligere hen imod røntgenbillederne med en forsimplet brud/ikke-brud beslutning. Det kræver med andre ord ekstra opmærksomhed og

konferering at bryde med fraktur-AI's forsimplinger. Rasmus' håndbevægelse og bemærkninger om at blive snydt og frygten for at blive "doven" peger også på denne mulighed for, at fraktur-AI kommer til at dominere processen. De kropslige, affektive og gennem konferering og oplæring kollegialt opnåede former for viden, som ligger i den kliniske situerede tid, hvor læger inddrager patientens fortid, nutid og fremtidsscener, overskygges derved af fraktur-AI's fastfrosne tid og binære afgørelse om brud/ikke-brud, som vi kalder en punktuering af tid. Når opmærksomheden trækkes mod AI'ens vurdering af røntgenbilledet som punkt, overskygges den fremtidige orientering, da fraktur-AI's kontekstløse markeringer ikke i sig selv fortæller noget om typen af brud eller behandlingsform. Da fraktur-AI ikke forholder sig til tidligere røntgenbilleder eller patientens historie, udelukkes også orienteringer mod fortiden. Vurderingen af patienten i skadestuen skubbes henimod en vurdering af røntgenbillederne, og vurderingen af røntgenbillederne skubbes henimod en fragmenteret, binær overvejelse om brud/ikke-brud. AI-markeringen fremlægges uden kontekst og kræver lægernes fortolkning og evner til at gennemskue og begrunde teknologiens resultater.<sup>5</sup>

Selvom røntgenbilleder i sig selv allerede punktuierer tid ved at give et øjebliksbillede, viser vores materiale, at fraktur-AI intensiverer denne proces. AI'en tilføjer en forhåndsfortolkning, som låser billedets betydning i en markering af brud, som fremstår autoritativ uden at være forankret i en klinisk kontekst. Dette reducerer de tidlige kompleksiteter, lægerne ellers inddrager gennem patienternes fortællinger, undersøgelser og løbende kliniske overvejelser. Fraktur-AI's automatiserede form for punktuering ændrer således forholdet mellem det punktuelle og det processuelle i klinisk beslutningstagning og skaber en særlig type tidslighed, som lægerne må arbejde aktivt på at håndtere eller ignorere. Flere erfarne læger havde konkrete eksempler på algoritmiske fejlorienteringer. En ortopædkirurgisk læge i skadestuen beskrev et specifikt eksempel med en patient, som havde fået en bold hårdt ind på sin tommelfinger.

*Manden havde ondt i et af tommelens led på den side, hvor der er nogle ledbånd, man skal passe på. På røntgenbilledet var der et lille hak i patientens knogle. Fraktur-AI markerede, at det var et brud, hvilket gjorde lægen forvirret, og som han forklarede i et interview, tænkte han først "uha, den stakkels mand har nok brækket sin tommel." Men da lægen for anden gang gik ind og undersøgte patienten, fik han øje på et stort ar hen over tommelfingeren, lige der hvor patienten havde ondt, og hvor "AI'en havde OBS'et for en fraktur." Lægen gennemgik gamle røntgenbilleder fra tidligere*

---

5. Vi kan ikke gå i dybden med denne pointe i denne artikel, men der er fare for, at AI ligesom andre digitale teknologier og infratrakturer gør fagfolks arbejde usynligt (jf. Star & Strauss, 1999; Justesen & Plesner, 2025).

*på året, som viste, at mandens tommelfinger tidligere havde været brækket på samme sted, hvor det nu gjorde ondt. Han vurderede, at patientens smerte var for lille til at indikere et nyt brud, og at det tidligere brud var fuldstændig helet. På baggrund heraf konkluderede lægen, at der ikke var behov for yderligere behandling i den aktuelle situation.*

Den ortopædkirurgiske læge beskriver i dette eksempel, hvordan de diagnostiske beslutninger i høj grad må tage højde for fortiden til at guide nutiden og fremtiden. Det ikke-normale "hak" på røntgenbilledet skal vurderes i forhold til patientens aktuelle smerter, og om dette derfor tilsammen kunne repræsentere et behandlingskrævende brud. Sådanne vurderinger kræver fortolkning og vægtning af informationer i den konkrete kontekst. De må forankres i både fortiden, nutiden og fremtiden og tydeliggør dermed de tidlige vekselvirkninger, der præger den diagnostiske proces.

Fraktur-AI understøttede ofte ikke de erfarne lægers vurderinger, men medførte derimod ekstra arbejde: De måtte forholde sig til et nyt sæt data fra AI'en og begrunde deres beslutninger i relation til algoritmens markeringer på røntgenbilledet. Som en ortopædkirurgisk læge beskrev det, følte det som at "slås med en algoritme". Vi foreslår, at denne forskydning i opmærksomheden, hvor fraktur-AI's fastfrosne tid og binære beslutning overskygger den diagnostiske proces ved at trække fokus mod et bestemt datapunkt, delvist forklarer, hvorfor flere erfarne læger oplevede teknologien som "irriterende", "distraherende" eller "frustrerende", mens nyuddannede læger i højere grad havde en anden opfattelse.

Vores analyse viser, at fraktur-AI's fastfrosne binære bud udgør en punktuering af tid, som får en særlig vægt, og som klinikerne må opholde sig særligt længe ved. Dette kan være til gavn f.eks. for yngre læger, som opnår en mulighed for konferering og sikkerhed i deres diagnosticering, men det kan også opleves som skade og irritation, enten fordi der ikke er brug for dette forstyrrende ophold eller fordi informationen om brud/ikke-brud er meget forenklet og ikke i sig selv fortæller særligt meget. Når fraktur-AI ikke i nævneværdig grad understøtter lægernes kliniske praksis, hvor diagnostik og behandling præges af vekslende og kontekstafhængige tidsligheder, bliver det relevant at undersøge, om teknologien i stedet understøtter andre områder. Det næste analytiske fokus retter sig derfor mod fraktur-AI's betydning i relation til en overordnet organisatorisk værdi om effektivitet forstået som tempo og minimering af ventetid.

## Tempo og ventetid

Med overskriften “Kunstig intelligens afgør røntgenbilleder og giver mindre kø på skadestuen” bekendtgjorde Region Nordjylland i juni 2023, at de havde implementeret fraktur-AI i regionens skadestuer og skadesklinikker (Region Nordjylland, 2023). Blot otte måneder senere blev implementeringen afrapporteret i regionale og nationale nyheder med overskrifter såsom “Kunstig intelligens på skadestuen: Ventetiden reduceret med 60 minutter” (Witten, 2024) og “Kunstig intelligens sænker ventetid og fejlrate på nordjyske skadestuer” (Redaktionen, 2024). Overskrifterne illustrerer, at “ventetid” var et centralt omdrejningspunkt for implementeringen af fraktur-AI, og at ventetid tager overskrifter. Den nationale mediedækning, som kom på bagkant af beslutningen i februar 2025 om at “opskalere” og udbrede fraktur-AI til hele landet, havde ligeledes fokus på ventetid (Nisgaard & Nielsen, 2025). I disse medieberetninger fremføres fraktur-AI som en teknologi, der kan understøtte effektiv beslutningstagning hos yngre, uerfarne læger, og som kan indfri de populære visioner om hurtigere, delvist automatiserede vurderinger og derved reducere ventetid grundet pressede eller manglende lægefaglige ressourcer.

I skadestuen refererer ventetid til *patientens* tid fra ankomstsregistrering til registrering af et afsluttet forløb (se figur 2). Det er et udtryk, som betoner patientens oplevede tid som stagneret og spild-af-tid. Endvidere udgør ventetid et tidsligt fokus på hurtighed som en central værdi for god behandling i sundhedssystemet og på tempo som et væsentligt parameter for at måle kvalitet. Dette har vi f.eks. i Danmark særligt set inden for kræftområdet, hvor hurtighed og stram tidsstyring fremstilles som et gode og et værn mod den potentielt dødelige ventetid (Tørring, 2023).

Figur 2 synliggør imidlertid nogle af alle de opgaver, som løses af forskellige læger og andre, mens patienterne opholder sig i venteværelset og “venter”. I dette afsnit viser vi, at de forskellige parter i projektet nok havde et fokus på ventetid og tempo, men ikke et direkte fokus på patienternes ventetid, hvilket havde været fokus i andre afprøvningsprojekter.

På projektmøder inden afprøvningsprojektet skulle gå i gang diskuterede ledere, klinikere, kliniske forskere og projektdeltagere fra Indkøb og Medikoteknik, hvilke effekter man kunne forvente ved brugen af fraktur-AI. Efter en række indledende møder blev det besluttet, at der ikke kunne afsættes de nødvendige omfattende ressourcer, som det ville kræve at foretage en grundig kvantitativ videnskabelig undersøgelse af AI-redskabernes kvalitet eller præcision. Det var til gengæld muligt at måle på udvalgte Key Performance Indicators (KPI'er), som

fokuserede på skadestuens produktion. Ventetid blev et af fire KPI'er, som den kvantitative del af projektet fremstillede data på. De tre andre var beskrivertid, det vil sige den tid det tager for radiologerne at beskrive et røntgenbillede, antal genindkaldelser af de enkelte patienter og antal supplerende undersøgelser bestilt efter røntgenundersøgelsen. Vægtningen af KPI'erne peger derved på et helt centralt styringsparameter for en velfungerende skadestue: tidsstyring.

I vores etnografiske materiale og interviews med læger var der ingen, der talte om ventetid som sådan. Til gengæld var fokus på "at sikre flowet [af patienter] i skadestuen". "Patientflow" er inspireret af management og logistik-metoden Real-Time Capacity Demand (RTCD) og fokuserer på patienternes vej igennem hele behandlingssystemet over tid (Dansk Selskab for Patientsikkerhed, 2015). At sikre patientflow i skadestuer indebærer f.eks. at sikre den tværfaglige koordinering mellem forskellige opgaver i skadestuen og i relation til andre afdelinger, så patienterne "flyder" igennem skadestuen mest muligt glidende og "afsluttes", når de er færdigbehandlede eller kan sendes videre til egen læge, andre afdelinger eller opfølgende ambulante konsultationer. Dette skal ske uden for mange flaskehalse, som skaber stagnation, og uden propper som direkte stopper passagen.

Over for os nævnte lægerne ofte patientflow som det, der karakteriserer godt skadestuearbejde, f.eks. denne yngre læge i en introduktionsstilling i skadestuen.

Interviewer: *"Hvad kan du godt lide ved akutmedicin?"*

Intro-læge: *"Patientflow. En masse forskellige patienttyper. Altså unge, gamle, mænd, kvinder. En varieret hverdag. [...] Det er der, hvor alt det, vi har lært og trænet, kan bruges rigtigt. Jeg kan godt lide teamsamarbejdet, når vi modtager patienter sammen, der er dårlige. [...] Jeg kan godt lide at se mange patienter på en dag. Jeg er ikke til lange, lange samtaler i et eller andet ambulatorie, hvor man sidder og har en time med hver patient."*

Et hurtigt og sikkert patientflow samt at få sendt patienter videre til de rigtige steder er en fælles opgave for skadestuen og noget, som læger med forskellige specialer og sygeplejersker arbejder sammen om. Patientflow for lægerne i skadestuen indikerede også et højt tempo med afvekslende opgaver og mange patienter igennem afdelingen på en arbejdsdag.

I bogen *Rationalizing Medical Work* (1997) viser Marc Berg, hvordan forskellige beslutningsstøtteknikker og -teknologier som kliniske protokoller eller klassiske computeralgoritmer siden 1950'erne i stigende grad er indgået i den videnskabelige og rationelle organisering af medicinsk praksis. Danske læger har tiltagende arbejdet inden for en logik om "evidensbaseret medicin" (Timmermans & Berg,

2003), som tilskynder standardisering og protokollering af medicinsk arbejde. Beslutningsstøtteredskeer er med til at cementere evidensbaseret medicinsk praksis og strukturere den som en række beslutningssekvenser, hvor det er lægens opgave at træffe de rigtige rationelle beslutninger, ofte baseret på statistisk beregnet evidens. Kropsligt forankrede fortolkninger af fysiske undersøgelser bliver til symptomer på diagnoser, f.eks. er trepunktsømhed et tegn på brud af scaphoideum, som behandles efter de gældende regimer. På vægge, laminerede ark og computerskærme i danske skadestuer finder man mange forskellige flowcharts, algoritmer og bogstavkombinationer, som minder personalet om forskellige protokoller, handlingsrækkefølger og guidelines til diagnosticering og behandling, og som af nutidens læger bruges aktivt og opleves som støtte i deres faglige arbejde.

Patientflow handler således om andet og mere end at "nedbringe unødigt ventetid", selvom dette også er en positiv gevinst (Dansk Selskab for Patientsikkerhed, 2015). Ligesom mediedækningen af fraktur-AI fokuserede leverandøren af et af AI-redskaberne dog særligt på patienternes ventetid som én af en række "observationspunkter" i forhold til at "kunne realisere en gevinst" ved brugen af fraktur-AI. På et fællesmøde med projektgrupperne fra de to hospitaler i september 2024 argumenterede leverandøren imod "en kultur, hvor man afventer lægen" og fremhævede i stedet værdien i "at få tømt venteværelset". I hele projektperioden argumenterede leverandøren for at ændre nuværende arbejdsgange, så de yngre læger og eventuelt sygeplejersker i skadestuen kunne undvære enten konferering med speciallæge eller de radiologiske beskrivelser af røntgenbillederne i selve situationen og i stedet selv at kunne tage beslutninger og tolke røntgenbillederne med resultaterne fra fraktur-AI og behandle patienten. Med andre ord, hvis skadestuens yngre læger kunne undvære enten konferering eller radiologiske røntgenbeskrivelser for en tid, så behøvede dette ikke stoppe arbejdsgangen i skadestuen. I leverandørens fremstilling af fraktur-AI blev redskabet en mulighed for at spare ressourcer (både tid og læger) i en klassisk logistik-optimerende tilgang.

Men som vores analyse af de diagnostiske og behandlingsmæssige tidsligheder i sidste afsnit viste, er praksis en anden i skadestuen. I praksis er de radiologiske beskrivelser som regel allerede flyttet både rumligt og tidsmæssigt væk fra arbejdsgangen i skadestuen, og ventetiden for patienterne er ikke relateret til de radiologiske beskrivelser, men påvirkes særligt af mængden af patienter i forhold til mængden af læger. Lægerne i skadestuerne afventer ikke radiologiske beskrivelser, og de er heller ikke blot interesserede i de radiologiske beskrivelser på grund af en binær usikkerhed om fraktur/ikke-fraktur, som AI'en viser. De yngre

læger har brug for lægefaglig sparring til at bestemme forskellige typer af brud og diagnostiske nuancer samt de bredere behandlingsmæssige konsekvenser.

Der var en udbredt opfattelse blandt både yngre og mere erfarne læger på tværs af specialer, at fraktur-AI i nogle tilfælde kunne give de yngre læger en form for øget tillid eller sikkerhed til deres egne vurderinger – særligt i de tilfælde, hvor hverken den yngre læge eller fraktur-AI havde fundet noget brud. Men vores interview viste, at det, som gav de yngre læger, som var alene på vagt i skadestuen om natten, den største sikkerhed, var bevidstheden om, at deres vurderinger og beslutninger næste morgen ville blive sammenholdt med radiologernes beskrivelser af røntgenbillederne af en erfaren læge til røntgenkonferencen. En yngre læge fortæller:

*Vi har en røntgenkonference, så det er røntgenlægerne, der beskriver røntgenbillederne. Og så sidder vores bagvagter, altså en af de gamle ortopædkirurger, og kigger på svaret og læser, hvad vi har skrevet, os der har set patienten, og så vurderer de, om de vil ringe eller sende et brev. [...] Så det er nogle af de rutinerede, der tager beslutninger, og det er jo for mange unge en god tryghed at vide, at der bliver kigget på det hele igen. Vi har generelt en tommelfingerregel, for eksempel om natten, så overbehandler vi en lille smule, hvis man er i tvivl, fordi man altid kan klippe en gips op.*

Denne røntgenkonference, og den kliniske erfaring og radiologiske ekspertise, som mødes dér, var der ingen, som stillede spørgsmålstejn ved, eller som nogen forventede kunne blive erstattet af AI.

Når yngre læger i skadestuen giver udtryk for, at de om natten, når de er alene på vagt, kan bruge fraktur-AI til at træffe hurtigere beslutninger – f.eks. om at lægge en gips for en sikkerheds skyld – skyldes det ikke fraktur-AI *i sig selv*. I stedet er det muligheden for, at de yngre læger kan vægte fraktur-AI som et argument i deres kliniske diagnosticerings- og beslutningsproces. Dette kan de gøre, fordi røntgenkonferencen den førstkommande hverdag udgør et sikkerhedsnet.

I dagstid er det heller ikke selve AI'ens markering om brud eller ikke-brud *i sig selv*, som kan øge tempo eller begrænse ventetid i skadestuen. Indimellem var der nogle, især behandlersygeplejersker<sup>6</sup> og teknisk orienterede projektdeltagere, der berørte muligheden for opgaveglidning, hvor behandlersygeplejersker under lægernes diagnosticerings- og behandlingsansvar kunne anvende AI til at vælge bestemte former for bandagering eller til at udskrive patienter. Dette ville kræve en ændring af de organisatoriske rammer og lokale instrukser i skadestuerne. Andre gange grinede lægerne højt og tegnede et absurd billede af fremtidens

---

6. En behandlersygeplejerske er en videreuddannet sygeplejerske med kompetencer og ansvar til at udføre visse former for lægelig behandling.

skadestuer, hvor “der er folk, som får et 14-dages kursus i at gipse, som klarer skadestuen – og så klarer computeren resten.” For de kliniske deltagere var det afgørende, at der blev skabt klinisk evidens for fraktur-AI, men dette ville kræve større kliniske studier end nærværende afprøvning af fraktur-AI, hvor ventetid gled i forgrunden foran et fagligt fokus på patientflow.

I nærværende artikels udforskning af, hvorfor der overhovedet var tilslutning i skadestuerne til at afprøve fraktur-AI, når praktikerne tydeligvis anså dets funktion som begrænset og kontekstafhængig, er det endnu en form for tidslighed, som tilbyder svar: De løfter, som AI kommer med, og de forestillinger om, hvor god og præcis fraktur-AI's påvisning af brud antages at blive i fremtiden.

## AI-som-løfte og forventninger til fremtiden

I vores undersøgelse oplevede mange læger, at fraktur-AI lavede forskellige fejl og oftest ikke fungerede som beslutningsstøtte. Alligevel havde de generelt en positiv indstilling til AI, hvilket forekom os at være et paradoks (jf. Hoeyer, 2023). Når vi nærlæste vores interview og noter fra feltarbejdet, tegnede der sig et billede af, at fraktur-AI *her og nu* ikke kunne bruges til meget (andet end til de yngre læger om natten, se ovenfor), men at mange fremhævede, at fraktur-AI, ligesom andre former for AI, i *fremtiden* ville være et rigtigt godt redskab til beslutningsstøtte. Berettigelsen af afprøvningen af fraktur-AI afhang altså af de løfter, som lægerne forbandt med AI. Selvom der ikke var god evidens for at bruge AI her og nu, herskede der ingen tvivl om, at “AI er kommet for at blive”, og at det var vigtigt for læger at lære at bruge AI. I dette afsnit viser og diskuterer vi dette paradoks og sammenligner det med det, som Klaus Høyer (2023) har skrevet om dataparadokser, nemlig at data-som-løfter har større indflydelse og betydning end data-som-evidens. Som tidslighed vægter AI-som-løfte således fremtiden og endnu ikke realiserede forventninger over håndgribelige, nutidige erfaringer.

Flere læger på tværs af erfaring og speciale talte om, at fraktur-AI i fremtiden ville være en stor hjælp og et sikkerhedsnet. Men disse forhåbninger blev som regel knyttet sammen med kvaliteter i fraktur-AI, som teknologien endnu ikke havde. En yngre læge fortalte i et interview:

Interviewer: *“Kan du se et behov for at bruge frakturalgoritmen? Synes du, der er behov for den i jeres arbejde?”*

KBU-læge: *“Ja, det tænker jeg da. Specielt når den bliver bedre. Jeg tænker også på fremtiden. Så er den så tæt på 100% rigtig, at det er da supergodt. Altså, så misser*

*man ikke ting. Specielt midt om natten, når man er alene og er en uerfaren KBU, at den så med sikkerhed kan sige, at det er et brud eller ej. Eller med meget højere sikkerhed, det er da meget godt. Lige nu er det jo lidt sådan, kan man sige, man skal være forsigtig med at stole på det, når den siger ikke-brud."*

Der er noget ironisk eller ligefrem paradoksalt ved, at denne læge lige nu er "forsigtig med at stole på" AI-resultaterne samtidigt med, at vedkommende var sikker på, at den i fremtiden ville blive "tæt på 100% rigtig". Mange andre læger havde ligeledes opfattelsen af, at fraktur-AI med tiden ville blive "fejlfri", "bulletproof" eller få "100% træfsikkerhed", som de udtrykte det. Disse forhåbninger peger på ideen om, at algoritmens performance potentielt kunne øges, og at AI desuden ville kunne udvikles til ikke kun at markere frakturer, men også at stille diagnoser om det helt præcise brud, måle vinkler på knoglestillinger og meget andet. Det var en udbredt forestilling, at fraktur-AI kunne "lære", hvormed der menes, at lægernes beslutningsinput ville gøre redskabet bedre til at afgøre, om der er fraktur eller ej, selvom ingen af de specifikke afprøvede AI-redskaber var dynamiske på denne måde og denne form for træning pt ikke er mulig, bl.a. på grund af juridiske beslutninger. Det er forfatterne bekendt ikke blevet nævnt i forbindelse med dette projekt ved AUH og RHH, så vi antager, at denne forestilling om "læring" kommer fra bredere samfundsmæssige fortællinger om AI og machine learning. Metaforen machine learning foregiver netop, at computere med AI-programmer bliver ved med at "lære" og forbedre sig selv, hvilket de i de færreste tilfælde kan. Det springende punkt for den positive opfattelse af fraktur-AI i nutiden var med andre ord påvirket af, at der i fremtiden ville blive nok evidens til, at der ville være det, som blev kaldt "rygdækning", til at bruge redskabet, fordi både yngre og erfarne læger ville kunne "stole på" redskabets resultater.

Et af de dataparadokser, som Høyer (2023) beskriver i sin monografi om sundhedsdata og deres politiske betydning, er at kravet om data som evidens for at træffe politiske, organisatoriske eller økonomiske beslutninger på den ene side fører til, at der skabes og registreres en stor mængde data til senere brug, men at disse data til gengæld sjældent lægges til grund for konkrete beslutninger i nutiden. Der er en stor diskrepans mellem, hvad data angiveligt skal bruges til i fremtiden, og hvad data rent faktisk bruges til i nutiden, og data forventes at levere fremtidig evidens for beslutninger, som allerede er truffet og derved ikke informerer beslutninger i nutiden (Høyer, 2023, s. 46, 52). På samme paradoksale måde vil vi sige, at det bedste argument for at bruge fraktur-AI nu er, at AI-teknologier vil være rigtig gode i fremtiden.

I forlængelse af dette kan vi trække på Stevens et al. (2018), som viser, hvordan

forståelser af big data i sundhedsvæsenet ofte diskursivt ledsages af særlige epistemologiske forventninger. Forfatterne viser, hvordan brugen af big data legitimeres gennem forestillinger om det fremtidige bidrag til en mere objektiv og præcis lægelig praksis. Data beskrives som noget, der blot findes “derude”, parat til at blive “høstet”. Disse løfter er performative, idet de skaber en forventning om forbedringer, som igen legitimerer yderligere investeringer og afprøvninger. På samme måde bliver afprøvningen af fraktur-AI i vores materiale ikke først og fremmest meningsfuld gennem redskabets aktuelle anvendelighed, men gennem de fremtidige kvaliteter og forbedringer, som lægerne forventer, at den vil indfri.

Som Watson og Wozniak-O'Connor (2025) fremhæver i deres analyse af “the promise of AI in health”, er det netop dette fremtidsorienterede løfte, der i stigende grad fungerer som en institutionaliseret ramme på sundhedsområdet. AI fremstilles ikke blot som et nyttigt redskab, men som en uundgåelig nødvendighed, hvis legitimitet hviler på forestillingen om, at teknologien snart vil levere på sine løfter. Derved udskydes mulig kritik af redskaberne i nutiden, da det jo er på vej til at blive bedre. Det er med andre ord løftets potentialer, som bliver meningsbærende. Dette svarer til, hvordan lægerne i vores undersøgelse taler om fraktur-AI: De oplever dens mangler i praksis, men fastholder alligevel forventningen om, at den vil blive “100% rigtig” i fremtiden.

Inden for videnskabs- og teknologistudier (STS) er der en lang tradition for at undersøge forskellige nye teknologiers “performativitet”, det vil sige den selvpfyldende profeti, der ligger i, at de bliver udråbt som fremtiden, og at der derfor på alle mulige måder investeres i dem. AI udgør i dag denne type løfterige og performative teknologi (Watson & Wozniak-O'Connor, 2025). Dette viser sig i en dansk sammenhæng i Regeringens strategier for kunstig intelligens (2019, 2024), senest forventningen om at kunne “frigør[e] mindst 50 mio. timer svarende til mindst 30.000 årsværk på tværs af den offentlige sektor frem mod 2035” (Digitaliseringsministeriet, 2025) og i de forventninger, der ligger i store offentlige organisationer såsom hospitaler, til, at man skal teste og implementere AI-teknologier på trods af stor usikkerhed om deres positive effekter. Ud fra et “forventningssociologisk” perspektiv (Brown & Michael, 2003) kan man sige, at projektet “afprøvning af fraktur-algoritme” indebar deltageres “håbefulde praktiseren (wishful enactments)” (Tutton, 2011) af AI som forskellige teknologier, der er “kommet for at blive”.

Arbejdet med fraktur-AI udgør med dette perspektiv tidslige “ønskehandling” eller “forventningshandling”, hvor deltagerne tilslutter sig ønsket og forventningen om, at AI vil blive til noget en dag, som det dog endnu ikke er og måske aldrig bliver. Det samme gjaldt en række andre former for AI, som lægerne

i skadestuen havde mere eller mindre erfaring med, og som de nævnte i interviews og samtaler.

Det interessante og paradoksale i vores materiale er, at disse forventninger til AI i fremtiden står i direkte modsætning til ærgrelserne med AI i hverdagen, som de mange udsagn om, at fraktur-AI "ingen værdi har", "ikke gør nogen forskel" eller "ikke er til hjælp", illustrerer – men at AI-som-løfte alligevel "uden tvivl" anses som realiserbart. Fælles var også en stemning af, at AI under alle omstændigheder var "godt at prøve af", for "der bliver kun mere AI" i sundhedssystemet i fremtiden. Til et af projektmøderne med de erfarne røntgen- og skadestuelæger konkluderede projektlederen, at projektets vigtigste formål "handler om kompetenceudvikling for klinikken, og hvordan klinikere modtager AI." Med dette blev det fremhævet, at det var vigtigt, at klinikere i praksis begyndte at få erfaringer med AI, og at skadestuen var et godt sted, fordi mange kommende læger kommer forbi skadestuen i deres KBU. Projektet kan her ses som en strategisk afprøvning, hvor formålet ikke først og fremmest var at måle eller beskrive det konkrete redskabs effekter, men "at modne" klinikken og organisationen og derved positionere den til fremtidige implementeringer af AI.

Ikke alle forventninger til AI var af positiv karakter. Der var også forventninger, som bar præg af bekymringer. En central bekymring handlede om lægernes evne til at kunne "tænke selv". På tværs af lægernes erfaringsniveau blev det fremhævet, at en kernekompetence for lægen er at tænke diagnostisk og behandlingsmæssigt på baggrund af at have indsamlet relevante informationer. I arbejdet med røntgenbilleder omhandlede dette "en kritisk gennemgang" og en "systematik" i at kigge hele billedet grundigt igennem. Når Rasmus, ligesom vi så det hos mange andre yngre læger, skærmede for fraktur-AI'ens svar med hånden, var det ud fra ønsket om at lære at læse røntgenbilleder og beholde denne kompetence som et nødvendigt led i at være en god læge. Med andre ord blev muligheden for at få hjælp til at vurdere et røntgenbillede også en risiko for at blive frataget muligheden for at lære "at stole på min egen vurdering", som en yngre læge i skadestuen forklarede. Den beslutningsstøtte, som i første omgang blev set som en positiv forventning til AI, blev derved også tolket ind i bekymringen om "at blive dårlige, at blive slappe", netop fordi AI-redskabet ville kunne gøre arbejdet i fremtiden og tilbyde en mulig genvej i deres beslutninger. På tværs af erfaringsniveau og speciale så vi blandt lægerne dette fokus på læring og bekymringer over at bruge et automatiseret hjælpemiddel. De positive forventninger fandt dog også vej i denne sammenhæng, idet den hurtige feedback, som AI-markeringerne giver, kunne blive værdifuld for lægernes læring. Til dette vil vi tilføje, at denne form for læring ved hjælp af AI vil kræve, at der vies tid og ressourcer til formålet, og

at der skabes de nødvendige læringsrum (Bruun et al., 2026). Opsummerende vil det sige, at der ligger en udfordring i, at læger og hospitaler på den ene side skal gå med tiden og bruge nye AI-redskaber, men at dette kræver ressourcer, som må fraregnes de forventede besparelser, hvis det ikke skal gå ud over personalets færdigheder til at arbejde med røntgenbilleder og andre teknologier.

Set i lyset af denne analyse kan vi forstå AI-som-løfte som mere end blot et lokalt paradoks i skadestuen. Det er en del af en bredere socioteknisk dynamik, hvor fremtiden fungerer som drivkraft for både politiske prioriteringer, organisatoriske investeringer og klinisk afprøvning, også selvom den nuværende teknologi fremstår mangelfuld.

## Konklusion

Som vi beskrev i indledningen af artiklen, forklarede den yngre, ortopædkirurgiske læge Fie under et interview, hvordan hun kort forholdt sig til, "at AI'en er der", for derefter hurtigt at vende tilbage til sin kliniske beslutningstagning. Hendes udsagn illustrerer et centralt træk ved vores analyse: At fraktur-AI ikke kan reduceres til et spørgsmål om diagnostisk præcision, men indgår i et socioteknisk system, som interagerer med kliniske praksisser, organisatoriske logikker og politiske forventninger.

I denne undersøgelse bliver fraktur-AI et eksempel på, hvordan AI-teknologier i første omgang viderefører logikker fra tidligere beslutningsstøteteknikker. Ligesom de ældre kliniske algoritmer og retningslinjer, baserer fraktur-AI sig på statistisk evidens og en statisk beregningsmodel, der skal leve op til krav om transparens og kvalitetsgodkendelse i sundhedsvæsenet. Teknologien er dermed ikke dynamisk lærende, men udtryk for en form for kodificeret klinisk viden, som indlejres i et digitalt format.

For det første har vi vist, hvordan fraktur-AI medfører en punktuering af tid, hvor diagnostiske processer fastfryses i en binær brud/ikke-brud-vurdering. Denne forenkling kan give lærings- og tryghedsmæssige gevinster for yngre læger, men overskygger de tidlige vekselvirkninger, som er centrale for erfarne klinikers situerede beslutningstagning. Teknologien kan således skabe nye forstyrrelser snarere end at understøtte den nuværende kliniske praksis.

For det andet har vi vist, hvordan tempo og ventetid fungerer som organisatoriske nøgleparametre, der legitimerer afprøvningen af fraktur-AI. Hvor leverandører og medier fremhæver teknologien som en løsning på patienternes ventetid, er det i praksis snarere skadestuens patientflow og tværfaglige samarbejde, der

former rytmerne. Fraktur-AI er her endnu ikke indlejret som et redskab, der reelt accelererer eller letter de kliniske arbejdsgange, men fungerer som et argument i diskussioner om effektivitet og ressourcenyttelse.

For det tredje har vi vist, hvordan fraktur-AI i høj grad bæres af AI-som-løfte. Selvom lægerne ofte oplever teknologien som fejlbehæftet eller direkte forstyrrende, fastholdes troen på, at den i fremtiden vil blive uundværlig og nærmest ufejlbarlig. Dette fremtidige blik bliver et performativt fundament, som legitimerer afprøvning og investering her og nu, også i fraværet af nutidig evidens for klinisk værdi.

Samlet set viser vores undersøgelse, at fraktur-AI ikke blot er et teknisk redskab til diagnostik, men et socioteknisk system, der både former og formes af skadestuens tidsligheder, organisatoriske logikker og politiske forventningshorisonter. Teknologien indgår i spændingsfeltet mellem klinisk praksis og politiske visioner, mellem nutidens irritation og fremtidens løfter.

Som socioteknisk system fremstår fraktur-AI stadig umodent, og de forventede gevinster i form af automatisering og effektivisering er organisatorisk diffuse. I vores undersøgelse flyttede vi fokus væk fra radiologernes arbejde og hen mod det tværfaglige samarbejde i skadestuen, hvor AI skal finde sin plads i en allerede presset og tidsfølsom praksis. Ligesom i studier af radiologer (Anichini et al., 2024; Lebovitz et al., 2022; Lombi & Rossero, 2024) viser vores materiale, at arbejdet med fraktur-AI giver anledning til friktion og kræver løbende fortolkning og tilpasning. I modsætning til radiologerne indgår skadestuens læger i et direkte, kropsligt og tidsligt møde med patienterne, hvor AI ikke blot påvirker diagnostiske vurderinger, men også rytmen i patientflowet og prioriteringen af opgaver. Fraktur-AI bliver dermed et redskab, der midlertidigt binder opmærksomheden, forskyder kliniske beslutningsrum og synliggør spændingen mellem ønsket om hurtigere behandling og behovet for faglig refleksion.

Sammenlignet med Marc Bergs (1997) klassiske analyser af beslutningsstøtte viser fraktur-AI dog en ny og mere begrænset form for rationalitet. Hvor tidligere redskaber som flowcharts og protokoller indlejrede beslutningstagning i en sekventiel proces, en form for hvis-så-system, der knyttede vurdering til næste trin og kontekst, så er fraktur-AI reduceret til et øjebliksbillede. Med sine markeringer på røntgenbilledet frembringer den en binær afgørelse: brud eller ikke-brud. Den tilbyder ingen proces, ingen kobling til klinisk ræsonnement og ingen tidslig forankring i patientforløbet. I den forstand udfordrer fraktur-AI selve definitionen af beslutningsstøtte, fordi kliniske beslutninger altid er situerede og afhænger af en forståelse af tid, rækkefølge og relationer.

Analysen peger således på, hvordan implementering af AI i sundhedsvæsenet ikke alene handler om effektivitet, men også om hvem og hvad der får autoritet til at definere klinisk praksis. Artiklen fremhæver dermed, hvordan divergerende interesser, værdier og forventninger bliver bærende i bestræbelserne på at føre AI-teknologier som fraktur-AI frem og gøre dem meningsfulde i en klinisk hverdag.

## Referencer

- Albinus, N.-B. (2022). *Mangel på specialiserede radiologer driver væksten hos TMC*. <https://dagensmedicin.dk/mangel-paa-specialiserede-radiologer-driver-væksten-hos-tmc/>
- Anichini, G., Natali, C., & Cabitza, F. (2024). Invisible to Machines: Designing AI that Supports Vision Work in Radiology. *Computer Supported Cooperative Work (CSCW)*, 33(4), 993-1036. <https://doi.org/10.1007/s10606-024-09491-0>
- Berg, M. (1997). *Rationalizing Medical Work: Decision-Support Techniques and Medical Practices*. MIT Press.
- Bowker, G. C. & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.
- Brown, N., & Michael, M. (2003). A Sociology of Expectations: Retrospecting Prospects and Prospecting Retrospects. *Technology Analysis & Strategic Management*, 15(1), 3-18. <https://doi.org/10.1080/0953732032000046024>
- Bruun, M. H., Antonsen, C. E., & Frumer, M. (2026, under udgivelse). The Role of AI in Clinical Learning: Revealing Tensions Between Patient Flow and Competence Development. I L. Smirnova & I. Engeness (red.), *Sociocultural Perspectives on AI-Enhanced Learning*. Springer.
- Dansk Selskab for Patientsikkerhed. (2015). *Sikkert Patientflow – Erfaringer fra et forbedringsprojekt*. <https://patientsikkerhed.dk/projekter/sikkert-patientflow/>
- Dansk Selskab for Patientsikkerhed. (2023). *Digital sundhedsfaglig behandling – pligter og ansvar. En guide til sundheds-faglige behandlingssteder*. <https://stps.dk/udgivelser/2023/okt/guide-om-digital-sundhedsfaglig-behandling>
- Danske Regioner. (2025). *40,6 mio. kr. til at brede kunstig intelligens-løsninger ud til flere dele af landet*. <https://www.regioner.dk/services/nyheder/2025/februar/40-6-mio-kr-til-at-brede-kunstig-intelligens-loesninger-ud-til-flere-dele-af-landet/>
- Digitaliseringsministeriet. (2025). *Mere tid til det vigtige. Målbillede for udrulning og anvendelse af kunstig intelligens i den offentlige sektor*. <https://www.digmin.dk/digitalisering/publikationer/2025/mere-tid-til-det-vigtige-maalbillede-for-udrulning-og-anvendelse-af-kunstig-intelligens-i-den-offentlige-sektor>
- Frumer, M. (2023). "Keeping an Eye on It": Infrastructures of Lung Cancer Uncertainty and Certainty. In R. S. Andersen & M. L. Tørring (Eds.), *Cancer Entangled: Anticipation, Acceleration, and the Danish State*. Rutgers University Press.
- Hammersley, M. & Atkinson, P. (2019). *Ethnography: Principles in Practice*. Routledge.
- Hoeyer, K. (2023). *Data Paradoxes: The Politics of Intensified Data Sourcing in Contemporary Healthcare*. MIT Press.
- Huisman, M., Ranschaert, E., Parker, W., et al. (2021). An international survey on AI in radiology in 1041 radiologists and radiology residents part 2: expectations, hurdles

- to implementation, and education. *European Radiology*, 31(11), 8797-8806. <https://doi.org/10.1007/s00330-021-07782-4>
- Justesen, L. & Plesner, U. (2025). *Digitalisering og usynligt arbejde*. Hans Reitzels Forlag.
- Katal, S., York, B., & Gholamrezaezhad, A. (2024). AI in radiology: From promise to practice – A guide to effective integration. *European Journal of Radiology*, 181, 111798. <https://doi.org/https://doi.org/10.1016/j.ejrad.2024.111798>
- Khalifa, M., & Albadawy, M. (2024). AI in diagnostic imaging: Revolutionising accuracy and efficiency. *Computer Methods and Programs in Biomedicine Update*, 5, 100146. <https://doi.org/https://doi.org/10.1016/j.cmpbup.2024.100146>
- Kidwai, A. S. (2023). *The new radiology landscape: Teleradiology and AI*. <https://www.solventum.com/en-us/home/health-information-technology/resources-education/blog/2023/5/the-new-radiology-landscape-teleradiology-and-ai/>
- Kristensen, D. B. (2025). *Professor: Brystkræftscreening med AI giver hurtigere resultater, men det øger risikoen for fejl*. <https://www.altinget.dk/digital/artikel/professor-brystkræftscreening-med-ai-giver-hurtigere-resultater-men-det-oeger-risikoen-for-fejl>
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science*, 33(1), 126-148. <https://doi.org/10.1287/orsc.2021.1549>
- Lombi, L., & Rossero, E. (2024). How artificial intelligence is reshaping the autonomy and boundary work of radiologists. A qualitative study. *Sociology of Health & Illness*, 46(2), 200-218. <https://doi.org/https://doi.org/10.1111/1467-9566.13702>
- McCullough, C. H., & Rajiah, P. S. (2023). Milestones in CT: Past, Present, and Future. *Radiology*, 309(1), e230803. <https://doi.org/10.1148/radiol.230803>
- Nisgaard, A., & Nielsen, S. A. (2025). *Kan du finde fejlen på billedet? AI kan - og det kan få betydning for din næste tur på skadestuen*. <https://www.dr.dk/nyheder/viden/teknologi/kan-du-finde-fejlen-paa-billedet-ai-kan-og-det-kan-faa-betydning-din-naeste>
- Olsen, T. L., Martensen, M. M., Yding, H., & Futtrup, L. R. (2023). *Minister vil bruge kunstig intelligens og rollebytte til at løse mangel på radiologer og radiografer*. <https://www.dr.dk/nyheder/indland/minister-vil-bruge-kunstig-intelligens-og-rollebytte-til-loese-mangel-paa-radiologer>
- Pesapane, F., Codari, M., & Sardanelli, F. (2018). Artificial intelligence in medical imaging: threat or opportunity? Radiologists again at the forefront of innovation in medicine. *European Radiology Experimental*, 2(1), 35. <https://doi.org/10.1186/s41747-018-0061-6>
- Pfaffenberger, B. (1992). Social Anthropology of Technology. *Annual Review of Anthropology*, Vol 21: 491-516. <https://doi.org/https://doi.org/10.1146/annurev.an.21.100192.002423>
- Redaktionen. (2024). *Kunstig intelligens sænker ventetid og fejlrate på nordjyske skadestuer*. <https://sundhedspolitisktidsskrift.dk/nyheder/behandlinger/8611-kunstig-intelligens-saenker-ventetid-og-fejlrate-pa-skadestuer.html>
- Regeringen. (2019). *National strategi for kunstig intelligens*. <https://regeringen.dk/nyheder/2019/national-strategi-for-kunstig-intelligens/>
- Regeringen. (2024). *Strategisk indsats for kunstig intelligens*. <https://www.digmin.dk/digitalisering/publikationer/2024/strategisk-indsats-for-kunstig-intelligens-et-styrket-fundament-for-ansvarlig-udvikling-og-anvendelse-af-kunstig-intelligens-i-danmark>
- Region Nordjylland. (2023). *Kunstig intelligens afgør røntgenbilleder og giver mindre kø på skadestuen*. <https://rn.dk/service/nyhedsbase-rn/2023/06/kunstig-intelligens-afgoer-roentgenbilleder-og-giver-mindre-koe-paa-skadestuen>
- Saunders, B. F. (2008). *CT Suite: The Work of Diagnosis in the Age of Noninvasive Cutting*. Duke University Press.

- Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*, 4(2), 2053951717738104. <https://doi.org/10.1177/2053951717738104>
- Star, S. L. (1999). The Ethnography of Infrastructure. *American Behavioral Scientist*, 43(3), 377-391. <https://doi.org/10.1177/00027649921955326>
- Star, S. L. & Strauss, A. (1999). Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work. *Computer Supported Cooperative Work (CSCW)*, 8(1-2), 9-30. <https://doi.org/10.1023/A:1008651105359>
- Stevens, M., Wehrens, R., & de Bont, A. (2018). Conceptualizations of Big Data and their epistemological claims in healthcare: A discourse analysis. *Big Data & Society*, 5(2). <https://doi.org/10.1177/2053951718816727>
- Suchman, L. (1987). *Plans and Situated Actions. The problem of human-machine communication*. Cambridge University Press.
- Timmermans, S., & Berg, M. (2003). *The gold standard: the challenge of evidence-based medicine and standardization in health care*. Temple University Press.
- Topol, Eric J. (2019). *Deep Medicine : How Artificial Intelligence Can Make Healthcare Human Again*. Basic Books.
- Tutton, R. (2011). Promising pessimism: Reading the futures to be avoided in biotech. *Social Studies of Science*, 41(3), 411-429. <https://doi.org/10.1177/0306312710397398>
- Tørring, M. L. (2023). The Waiting Time Paradox: Intensifying Public Discourses on the Vital Character of Cancer Waiting Times. In R. S. Andersen & M. L. Tørring (Eds.), *Cancer Entangled: Anticipation, Acceleration, and the Danish State* (pp. 23-41). Rutgers University Press.
- US Food & Drug Administration. (2025). *Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices*. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>
- Watson, A., & Wozniak-O'Connor, V. (2025). The promise of artificial intelligence in health: Portrayals of emerging healthcare technologies. *Sociology of Health & Illness*, 47(1), e13840. <https://doi.org/https://doi.org/10.1111/1467-9566.13840>
- Witten, P. (2024). *Kunstig intelligens på skadestuen: Ventetiden reduceret med 60 minutter*. <https://nordjyske.dk/nyheder/nordjylland/har-du-braekket-benet-saadan-har-skadestuen-mindsket-ventetiden/4792228>

# Personlige prognoser

## Forudsigelser af fremtiden med kunstig intelligens (AI) i klinisk praksis

Iben Mundbjerg Gjødsbøl<sup>1</sup>  
Mette Nordahl Svendsen<sup>1</sup>

---

<sup>1</sup>Center for Medicinske Videnskabs- og Teknologistudier, Institut for Folkesundhedsvidenskab, Københavns Universitet

[ibgj@sund.ku.dk](mailto:ibgj@sund.ku.dk)

[mesv@sund.ku.dk](mailto:mesv@sund.ku.dk)

Gjødsbøl, Iben Mundbjerg & Svendsen, Mette Nordahl. Personlige prognoser: Forudsigelser af fremtiden med kunstig intelligens (AI) i klinisk praksis. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 76-99. DOI: 10.7146/TFSS.V25I44.156294

Indsendt 05/25, accepteret 04/26, udgivet 06/26

I politik, sundhedsvæsen og forskning er der store forhåbninger til, at prædiktive algoritmer udviklet med kunstig intelligens (AI) vil individualisere og effektivisere forebyggelse, diagnostik og behandling af sygdomme. Forventningen er, at algoritmer, der kan integrere og processere store mængder og forskellige typer af patientspecifikke data, kan hjælpe sundhedsprofessionelle med at forudsige fremtidige hændelser såsom sygdom, komplikationer ved sygdom samt død for den enkelte patient for at kunne foregribe disse hændelser. Med andre ord fremstilles prædiktive algoritmer som nutidens orakler, der kan se ind i fremtiden med hidtil uset præcision ved at levere 'personlige prognoser' for borgere og patienter. I 2023 blev CARDIA<sup>IHD</sup>-algoritmen implementeret på forsøgsbasis i det elektroniske patientjournalssystem Sundhedsplatformen i Region Hovedstaden og Region Sjælland. Algoritmen prædikterer risikoen for at dø – eller prognosen for at overleve – for patienter, der indlægges med akut iskæmisk hjertesygdom på regionernes hospitaler. Baseret på et etnografisk feltarbejde blandt kardiologer, der behandler hjerte-

patienter, undersøger artiklen, hvordan kardiologer forstår og integrerer AI-drevne forudsigelser i klinisk praksis. Gennem empiriske eksempler viser vi, hvordan algoritmer som CARDIA<sup>IHD</sup> ikke kun leverer mere præcise forudsigelser af patienternes fremtider og ej heller blot fratager klinikere opgaver med at integrere og fortolke data. Derimod introducerer algoritmen nye former for epistemisk usikkerhed, som kræver, at klinikerne engagerer sig i nye former for fortolkningsmæssigt arbejde. Derudover kan klinikere have udfordringer med at afkode til hvad og hvordan, de kan bruge algoritmes prædiktioner i den kliniske beslutningstagning. Endelig viser vores analyse, at algoritmers tidlige horisonter ikke nødvendigvis er i overensstemmelse med den tids- og rumlige organisering af kliniske praksisser og patienternes forløb, hvilket udfordrer prædiktionernes relevans for den kliniske beslutningstagning.

## Personalized predictions: Predictive Artificial Intelligence (AI) in clinical practice

*In politics, healthcare, and biomedical research, there are high expectations that predictive algorithms developed with artificial intelligence (AI) will individualize and enhance the efficiency of prevention, diagnosis, and treatment. The hope is that algorithms capable of integrating and processing large volumes and diverse types of patient-specific data can assist health professionals in predicting future events such as disease, complications, and death for individual patients. In other words, predictive algorithms are portrayed as contemporary oracles that can anticipate the future with unprecedented precision by delivering ‘personalized prognoses’ for citizens and patients. In 2023, the CARDIA<sup>IHD</sup> algorithm was implemented on a trial basis in the electronic health record system ‘Sundhedsplatformen’ in Denmark’s Capital Region and Region Zealand. The algorithm predicts the risk of death – or the prognosis of survival – for patients admitted with acute ischemic heart disease at the regions’ hospitals. Drawing upon ethnographic fieldwork among cardiologists treating heart patients, this article examines how cardiologists understand and integrate the algorithm’s predictions into clinical practice. Through empirical examples, we show that predictive algorithms like CARDIA<sup>IHD</sup> do not necessarily deliver more precise predictions of patients’ futures, nor do they merely relieve clinicians of tasks related to integrating and interpreting data. Rather, the algorithm introduces new forms of epistemic uncertainty that require clinicians to engage in interpretive work. Moreover, clinicians may face challenges in decoding what the algorithm’s predictions can be used for and how they can be applied in clinical decision-making. Finally, our analysis demonstrates that algorithmic predictions do not necessarily align with the temporal and spatial organization of clinical practices and patients’ trajectories, which challenges the relevance of algorithmic predictions for clinical decision-making.*

## Indledning: Kunstig Intelligens (AI) og præcisionsmedicin

I politik, sundhedsvæsen og forskning udråbes Kunstig Intelligens (AI) i stigende grad som en central løsning på det dobbeltdemografiske pres, at et faldende antal sundhedsprofessionelle skal diagnosticere, behandle og drage omsorg for en stadigt større aldrende befolkning (Bajwa et al., 2021). AI og datadrevne teknologier forventes blandt andet at føre til hurtigere og mere præcis diagnostik, individuelt skræddersyet behandling af højere kvalitet og frigørelse af menneskelige ressourcer gennem automatisering og effektivisering af arbejdsgange (Johnson et al., 2021; Yu et al., 2018; Davenport & Kalakota, 2019). Politisk og forskningsmæssigt investeres der derfor massivt i metoder og værktøjer, der går under betegnelsen 'AI', og i både Robusthedskommissionen (2023) og Sundhedsstrukturkommissionens (2024) anbefalinger, samt i uddelingen af forskningsmidler, står en øget udvikling og brug af AI i sundhedsvæsenet helt centralt. I 2024 blev supercomputeren Gefion og det privat-offentlige firma Danish Centre for AI Innovation A/S indviet med en donation fra Novo Nordisk Fonden på 600 millioner kroner. Politisk forventes satsningen på og udviklingen af AI-drevne løsninger at skabe både bedre og sundere liv for borgere og patienter, et bæredygtigt sundhedsvæsen og en større samfundsmæssig rigdom.

I den kliniske praksis anvendes AI i stigende grad til at prædikere helbredsmæssige og medicinske risici (Haug & Drazen, 2023). Mulighederne for at integrere og processere hidtil usete mængder og typer af data, som AI tilbyder, forventes at føre til risikoprædiktioner, der er specifikke for den enkelte patient og dennes særlige karakteristika. Brugen af AI knyttes derfor tæt til ambitionen om at målrette forebyggelse, diagnostik og behandling til det enkelte individ ved brug af store mængder data. Denne tilgang kaldes personlig medicin eller præcisionsmedicin (Johnson et al., 2021; Prainsack, 2017). Potentialet i AI-drevne modeller, der bliver trænet på stigende mængder og modaliteter af populationsdata, er, at de kan integrere og finde mønstre i data på måder, der overgår den menneskelige regnekraft (Obermeyer & Lee, 2017; Topol, 2019). I de offentlige diskussioner fremstilles disse algoritmer og AI mere generelt som vor tids personlige 'orakler' (Frank, 2024), der kan forudsige fremtidige udfald med større præcision (Esposito et al., 2024; Dixon et al., 2024; Ali et al., 2023). Hos både politikere, forskningsfonde og biomedicinske forskere er håbet, at de 'personlige prognoser', som algoritmerne kan beregne, kan hjælpe de sundhedsprofessionelle med at forudsige den enkelte patients risiko for at blive syg, få komplikationer eller dø i fremtiden med henblik på at iværksætte forebyggende tiltag.

Ønsket om at forudsige sygdomme og sygdomsforløb er langt fra nyt i det medicinske felt, og de nye AI-drevne prædiktive værktøjer skriver sig ind i en lang række af teknologier, der beregner risici med henblik på at gøre fremtiden til genstand for intervention (Beck, 1992). Fra videnskabs- og teknologistudier (STS) ved vi imidlertid, at når nye teknologier bliver implementeret i praksis, forhandles deres brug og formål på andre måder end dem, der var intenderet eller forsøgt indlejret i teknologiernes 'scripts' (Akrich, 1992; Timmermans, 1996; Berg 1997). Fra et STS-perspektiv betragtes en algoritmes prædiktions dermed ikke blot som en beregnet sandsynlighed eller et udsagn om fremtiden, men som et fænomen, der bliver 'enacted' på mange forskellige måder i praksis, og som derfor er multipelt og situeret i tid og sted (Mol, 2002). I denne artikel undersøger vi, hvordan sundhedsprofessionelle forstår og integrerer AI-drevne forudsigelser som en ny vidensform i klinisk praksis. Empirisk tager vi afsæt i CARDIA<sup>IHD</sup>-algoritmen, der i 2023 blev introduceret på hospitaler i Region Hovedstaden og Region Sjælland som et risikovurderingsværktøj til kardiologer i et randomiseret klinisk forsøg. Algoritmen prædikerer overlevelsesprognosen for patienter, der indlægges på regionernes hospitaler med akut iskæmisk hjertesygdom, oftest blodprop i hjertet.

Artiklen er opbygget som følger: Først introducerer vi CARDIA<sup>IHD</sup>-algoritmen og den måde, den præsenteres på for kardiologerne i den elektroniske patientjournal, Sundhedsplatformen. Dernæst gør vi rede for de teoretiske begreber, der har informeret analysen, og herefter for de metoder og det empiriske grundlag, som artiklen bygger på. I den efterfølgende analyse demonstrerer vi gennem empiriske eksempler, at algoritmens prædiktions ikke nødvendigvis fratager klinikerne arbejdet med at integrere og fortolke data om den enkelte patient. I den praktiske brug introducerer algoritmen derimod nye former for fortolkningsmæssigt arbejde og epistemisk usikkerhed for klinikerne, og det viser sig vanskeligt for dem at anvende algoritmens prædiktions inden for de rammer, der organiserer det kliniske arbejde. Afslutningsvist diskuterer vi, hvordan vores analytiske fund nuancerer faglige og politiske idealer om transparens, explainability og det kliniske arbejdes organisering i samspillet mellem mennesker og algoritmer.

## CARDIA<sup>IHD</sup>-algoritmen

CARDIA<sup>IHD</sup>-algoritmen prædikerer chancen for at være i live efter et år for patienter, der bliver indlagt på hospitalet med akut iskæmisk hjertesygdom. Algoritmen er udviklet i et nordisk forskningsprojekt, Predict-IHD, med del-

tagelse af kliniske og bioinformatiske forskere fra Danmark, Norge og Island. Det overordnede formål med Predict-IHD-projektet var at udvikle et risikovurderingsværktøj, der beregner en mere præcis og individualiseret mortalitetsrisiko for patienter med akut iskæmisk hjertesygdom end eksisterende værktøjer ved at integrere flere data med betydning for patienternes prognose. Algoritmen er udviklet og trænet på sundhedsdata fra næsten 40.000 patienter med iskæmisk hjertesygdom trukket fra danske nationale registre og elektroniske patientjournaler. Disse data omfatter blandt andet diagnose- og procedurekoder, laboratorieresultater, koronararteriografier (KAG; røntgenundersøgelse af blodgennemstrømningen i hjertets kranspulsårer og dermed blodtilførslen til hjertet), medicin og kliniske målinger, og den endelige algoritme indeholdt op til 584 forskellige features. Algoritmen beregner automatisk en patients overlevelsesprognose cirka ti minutter efter, at patienten har fået lavet en koronararteriografi ved gennem Sundhedsplatformen at trække de patient-specifikke data, som prædiktionen beregnes ud fra. Efter udvikling blev CARDIA<sup>IHD</sup>-algoritmen valideret i en islandsk kohorte på 8.287 patienter. Det er videnskabeligt dokumenteret, at algoritmen demonstrerer forbedret nøjagtighed i forudsigelsen af overlevelsesprognosen for patienter med akut iskæmisk hjertesygdom sammenlignet med GRACE 2.0-risikoscoren (ikke refereret her af hensyn til anonymitet), som er det mest anerkendte værktøj til at estimere risikoen for at dø af alle årsager hos denne patientgruppe (D'Ascenzo et al., 2012).

Hypotesen, der har drevet udviklingen af algoritmen, er, at kardiologer, som kender den AI-beregne overlevelsesprognose for den individuelle patient, vil træffe bedre valg om behandling og opfølgning. I efteråret 2023 blev algoritmen integreret i Sundhedsplatformen som led i at afprøve algoritmen i et lodtrækningsforsøg (et randomiseret kontrolleret studie (RCT)). Sundhedsplatformen randomiserer automatisk patienter til enten interventionsgruppen (prædiktionen vises i Sundhedsplatformen) eller kontrolgruppen (prædiktionen beregnes men vises ikke i Sundhedsplatformen). Endepunkterne for dette studie er, om den kliniske brug af algoritmen vil forbedre et-års-overlevelsen for patienter i interventionsgruppen, og/eller minimere deres risiko for genindlæggelse sammenlignet med patienterne i kontrolgruppen. Algoritmen er altså implementeret i Sundhedsplatformen som et klinisk forsøg, der skal undersøge dens kliniske effekt, frem for blot at implementere den ukritisk. Spørgsmålet om, hvorvidt det vil forbedre patientens forløb, at de behandlende kardiologer kender CARDIA<sup>IHD</sup>-algorithms prædiktion, kan derfor først besvares, når resultaterne fra RCT-studiet er klar.

I Sundhedsplatformens datavisning bliver patienten placeret i en af i alt fem risikoklasser sat i forhold til den gennemsnitlige overlevelsesprognose beregnet ud fra en population af 12.000 patienter med iskæmisk hjertesygdom. De fem forskellige risikoklasser er visualiseret som en farveskala fra 'bedste 1-års overlevelse i forhold til gennemsnittet' (grøn) til 'lavere 1-års overlevelse i forhold til gennemsnittet' (rød). Udover selve prædiktionen præsenterer datavisningen også op til ti såkaldte 'forklarende faktorer', der skal hjælpe klinikerne med at forstå, hvorfor algoritmen prædikerer, som den gør. Sammen med prædiktionen præsenteres således fem faktorer, der driver prædiktionen mod død, og fem faktorer, der driver prædiktionen mod overlevelse med henblik på at øge algoritmens 'explainability'. Det er vigtigt at understrege, at de forklarende faktorer er matematisk korrelerede men ikke nødvendigvis kausalt forbundne med overlevelsesprognosen, og at hverken prædiktionen eller de forklarende faktorer giver konkrete anbefalinger til, hvilke beslutninger kardiologen skal træffe.

Det er således helt op til kardiologen selv, hvordan og hvorvidt algoritmens prædiktion bruges i den kliniske beslutningstagning. Uanfægtet algoritmens prædiktion skal kardiologer følge anbefalinger for procedurer og behandlinger med dokumenteret klinisk effekt; de såkaldte Klasse I anbefalinger i nationale og internationale guidelines (Eftekhari et al., 2025; Knuuti et al., 2020). Der er imidlertid også anbefalinger til behandlinger, for hvilke effekten er mindre veldokumenteret, og klinikere kan derfor overveje at justere disse Klasse II og III anbefalinger i lyset af algoritmens prognose. Eksempler på sådanne justeringer inkluderer mængden af opfølgende undersøgelser, antal, mængde og dosering af medicin, indlæggelsesvarighed, anbefalinger for rehabilitering og hyppigheden og intensiteten af den opfølgende behandling. På trods af at algoritmen ikke er eksplicit handlingsanvisende, er det altså hensigten, at dens prognoser, som tidsmæssigt rækker ud over de akutte risici, skal give anledning til overvejelser hos kardiologen om, hvor aggressivt og intensivt en given patient behandles i lyset af dennes langtidsprognose. Det betyder, at kardiologerne i hvert enkelt tilfælde selv skal forhandle, om algoritmen får karakter af beslutningsstøtte.

Derudover er der nogle særlige forudsætninger for CARDIA<sup>IHD</sup>-algoritmens prædiktioner, herunder at algoritmens prognose gælder i det øjeblik, den bliver beregnet, og at algoritmen ikke kan tage højde for fremtidige hændelser. Derudover skal klinikerne og sundhedsvæsenet mere generelt have sikret sig og registreret alle relevante parametre for en patient, så de er tilgængelige for algoritmen. Hvis en patient for eksempel udvikler hjertestop i forbindelse med en blodprop i hjertet og gennemgår langvarig genoplivning, der medfører mulig hjerneskade – og hjertestoppet ikke bliver registreret som et sådant (det vil

sige med en diagnosekode for hjertestop) i det akutte forløb – kan det påvirke algoritmens beregning i betydelig grad. Disse forhold er en påmindelse om, at en algoritmes prædiktioner afhænger af de data, som den er trænet på, og de data om den enkelte patient, som algoritmen har til rådighed for sin beregning.

På trods af de politiske og faglige forventninger til brugen af AI-drevne løsninger i sundhedsvæsenet og en stigende udvikling af prædiktive algoritmer, er det kun et fåtal af disse, der bliver implementeret og brugt i praksis – selv i et højt digitaliseret samfund og sundhedsvæsen som det danske. Set i lyset af de mange tekniske, juridiske og praktiske udfordringer, der er forbundet med at oversætte en kompleks matematisk model til klinisk brug, er det i sig selv exceptionelt, at det er lykkedes at integrere CARDIA<sup>IHD</sup>-algoritmen i Sundhedsplatformen, og at Sundhedsplatformen automatisk trækker og integrerer data samt randomiserer patienter til det kliniske lodtrækningsforsøg.

Baseret på samtaler med forskere i Predict-IHD-projektet og skriftligt materiale om algoritmen, som ikke er refereret her af hensyn til anonymitet, er det vores forståelse, at algoritmen går under betegnelsen 'AI', fordi den er udviklet som et neuralt netværk, der gør det muligt at integrere et bredere panel af features fra for eksempel patientjournaler og nationale registre end de allerede kendte risikoestimeringsværktøjer. De traditionelle værktøjer er baseret på klassiske statistiske teknikker og omfatter en begrænset mængde kendte risikofaktorer for iskæmisk hjertesygdom, der indtastes manuelt. Eksempler på sådanne risikovurderingsværktøjer er den førnævnte GRACE 2.0 risiko-score, HEART-scoren som forudsiger 30-dages risikoen for alvorlige kardiovaskulære hændelser (død, myokardieinfarkt eller akut revaskularisering) hos patienter med mistanke om akut koronart syndrom (Backus et al., 2010), og SCORE2, der prædikterer en persons samlede risiko for at udvikle hjertekarsygdom (SCORE2 working group & ESC Cardiovascular risk collaboration, 2021).

## Fremtid, risiko og algoritmiske prædiktioner

Teoretisk betragter vi prognose og tid mere generelt som socialt konstitueret og som noget, der etableres i praksis (Gjødsvøl & Svendsen, 2019; Willerslev et al., 2013; Mattingly, 1998; Fabian, 1983). I bogen *Futures Past: On the Semantics of Historical Time* argumenterer historikeren Reinhart Koselleck (1985) for, at et ændret forhold mellem erfaring og forventning markerede fremkomsten af moderniteten. Hvor mennesket i middelalderkirkens eskatologiske tænkning forholdt sig til fremtiden gennem apokalyptiske profetier, så forstår og former det moderne menneske sin fremtid gennem prognoser. Koselleck skriver om fænomenet prognose:

*"...time continually emanates from the prognosis in an unforeseeable, but predictable manner. Prognosis produces the time within which and out of which it weaves, whereas apocalyptic prophecy destroys time through its fixation on the End. From the point of view of prophecy, events are merely symbols of that which is already known"* (Koselleck, 1985, s. 19).

Hvor profetien repræsenterer en forståelse af tid, som altid er rettet mod verdens ende, og hvor fremtiden dermed fremstår kendt og lukket, så etablerer prognosen fremtidige begivenheder som både "unforeseeable" men samtidig "predictable". I den prognostiske tænkning er fremtiden kendetegnet ved ubestemmelighed og åbenhed, fordi prognoser skaber et rum for menneskelig intervention i nuet, der kan række ind i og forandre fremtiden (Koselleck, 1985).

I medicinsk praksis er beregning af risiko og sandsynlighed for fremtidige hændelser som for eksempel sygdom indbegrebet af den prognostiske tænkning. I en dansk kontekst har eksempelvis Schwennesen et al. (2011) vist og kritiseret, hvordan beregning af risici for Downs syndrom hos fostre er blevet en måde, hvorpå sundhedsvæsenet håndterer og intervenserer i en åben og usikker fremtid. I denne proces etableres fostret som risikoobjekt samtidig med, at forældrene forventes at agere som autonome og ansvarlige beslutningstagere, der handler rationelt på den foreliggende viden for at forhindre bestemte fremtider i at materialisere sig (Schwennesen et al., 2011). Denne form for 'risikologik' bliver kun yderligere forstærket af data-intensive, prædiktive algoritmer, som i stigende grad udvikles til brug i både sundhedsvæsenet og andre velfærdsstatslige institutioner med henblik på rettidig indgriben og forebyggende tiltag. Men hvor de forhenværende statistiske metoder beregnede sandsynlige fremtider på befolknings- og gruppeniveau, muliggør de nye prædiktive værktøjer risiko-beregninger, der er målrettet individet (Amoore, 2013). Dette forhold forandrer betydningen af usikkerhed i forhold til den måde, vi forholder os til fremtiden i hvert enkelt tilfælde. Således argumenterer filosofen Elena Esposito (2024) for, at algoritmiske forudsigelser indsnævrer den åbne fremtid, når mønstre, der er fundet i historiske data, bruges til at forudsige fremtidige begivenheder for individet:

*"Such [algorithmic] predictions [...] do not refer (as probabilistic calculations do) to the present observation of the future; rather they refer directly to the structures of the future itself, as they can be inferred from present data"* (Esposito et al., 2024, s. 5).

Værdien af den individualiserede prædiktion, der betragtes som grundpille i præcisionsmedicin, er altså, at den muliggør direkte indgriben i og formning af 'the structures of the future' for den enkelte patient. Det er imidlertid ikke så

meget korrektheden af forudsigelserne af fremtiden, der har værdi, men snarere det faktum, at prædiktionen bliver et redskab i nuet til at håndtere usikkerhed om fremtiden. I den medicinske praksis betyder det, at algoritmiske prædiktioner forventes at være brugbare i den kliniske beslutningstagning ved at bidrage med en retning – ‘an orientation’ (Esposito et al., 2024) – for den kliniker, der skal træffe beslutninger om, hvordan den individuelle patient bedst behandles.

Hvorvidt algoritmer får denne kraft til at forme fremtiden i praksis må undersøges lokalt, kontekstuel og empirisk frem for abstrakt (Ratner og Elmholdt, 2023). Inspireret af klassiske STS-begreber som ‘enactment’ (Mol 2002) og teknologisk indlejrede ‘scripts’ (Akrich, 1992; Timmermans, 1996) undersøger og analyserer vi i det efterfølgende empirisk, hvordan prognose aktivt etableres og forhandles i hverdagspraksis i kraft af relationer og interaktioner mellem klinikere, patienter og CARDIA<sup>IHD</sup>-algoritmen. Hvor algoritmiske prædiktioner forventes at reducere visse former for usikkerhed om fremtiden og dermed give en retning for de kliniske beslutninger, viser vores empiriske studie, at CARDIA<sup>IHD</sup>-algoritmen introducerer nye former for usikkerhed, og at klinikerne oplever udfordringer med at forstå og bruge dens forudsigelser i deres beslutningstagning (se også Tyskbo & Nygren, 2024; Lebovitz et al., 2022). Derudover kan algoritmens tidslighed ikke nødvendigvis integreres i de kliniske beslutninger, som allerede er formet af den måde, hvorpå patienters forløb er organiseret i tid og rum. Her trækker vi på Sharon Kaufmans (2006) begreb ‘pathways’, der hjælper os med at belyse, hvordan de tidlige og stedlige rammer for klinisk beslutningstagning aktualiserer bestemte fremtider, mens andre – heriblandt patientens AI-genererede overlevelsesprognose – nemmere udgrænses.

## Metode

Artiklen trækker på et etnografisk feltarbejde, som sigtede på at følge CARDIA<sup>IHD</sup>-algoritmen fra udvikling til implementering i klinikken (Amelang & Bauer, 2019). Feltarbejdet blev udført i perioden 2019 til 2024 af artiklens forfattere, der som samfundsvidenskabelige forskere i Predict-IHD-projektet havde til opgave at undersøge de sociale og etiske implikationer af at udvikle CARDIA<sup>IHD</sup>-algoritmen med særligt fokus på dens introduktion og brug i klinikken. Tilsammen har vi deltaget i 16 Predict-IHD-møder af én til halvanden times varighed. De første møder var fysiske møder, men grundet COVID-19 overgik møderne til online format. Som både deltagere og observatører på disse møder fik vi løbende viden om, hvordan arbejdet med at udvikle algoritmen skred frem. Udover

mødedeltagelse har vi foretaget ni kvalitative interviews med forskere fra Predict-IHD-projektet, der varede mellem én og halvanden time, og haft løbende udvekslinger med de danske forskere. I efteråret 2023 blev algoritmen integreret i Sundhedsplatformen på hospitaler i Region Hovedstaden og Region Sjælland. I ugerne op til algoritmens kliniske introduktion arrangerede PI for Predict-IHD-projektet en række præsentationer på hospitaler i Østdanmark, som behandler patienter med akut iskæmisk hjertesygdom. Her introducerede PI både det overordnede Predict-IHD-projekt og CARDIA<sup>IHD</sup>-algoritmen med et slideshow for de kardiologer, der var på arbejde den pågældende dag, enten før eller efter den daglige morgenkonference for at forberede hospitalspersonalet på at skulle bruge algoritmen i praksis. Disse præsentationer varede mellem 30 og 45 minutter og blev afsluttet med 15 minutters Q&A og fælles diskussion, hvor kardiologerne fik mulighed for at stille opklarende spørgsmål. Iben deltog i seks ud af i alt ti præsentationer, hvor hun udelukkende observerede introduktionen af algoritmen og den efterfølgende diskussion.

Endelig trækker artiklen på et klinisk feltarbejde, Iben udførte i efteråret 2023 og efteråret 2024, primært på et højt specialiseret hospital i København men med enkelte besøg til et andet hospital i København og et provinshospital. I alt fulgte Iben fem forskellige kardiologer men langt overvejende kardiologen Hans på det højt specialiserede hospital. Samlet set udførte Iben cirka 55 timers observation fordelt over otte dage, hvor hun fulgte kardiologer i deres arbejde med at behandle patienter med akut iskæmisk hjertesygdom. Formålet med dette feltarbejde var – i regi af Predict-IHD-projektet – at undersøge og få indsigt i, hvordan kardiologer bruger (eller undlader at bruge) CARDIA<sup>IHD</sup>-algoritmen i deres behandling af hjertepatienter, og hvordan de erfarer algoritmen som klinisk værktøj i deres beslutningstagning. Vores etnografiske positionering var således at undersøge den kliniske brug af algoritmen i hverdagspraksis uden et formål om at lade de etnografiske indsigter informere en eventuel videreudvikling af algoritmen.

I kraft af sin deltagelse i forskningsprojektet havde Iben umiddelbar adgang til at følge det kliniske arbejde, og hun fulgte ikke kardiologer, der selv havde været aktivt involverede i at udvikle algoritmen. Under feltarbejdet var Iben særligt opmærksom på, hvornår kardiologerne konsulterede algoritmen (eller undlod at gøre det), hvordan de forstod og fortolkede både prædiktionen og de tilhørende forklarende faktorer, og endelig hvordan de integrerede dens prædiktioner i den kliniske beslutningstagning. Analytisk har vi været inspireret af den abduktive tilgang (Tavory & Timmermans, 2014), der er karakteriseret ved en løbende vekselvirkning mellem teoretisk informeret analyse og feltarbejde med henblik på at nuancere forståelsen af, hvordan algoritmiske prædiktioner

integreres og mobiliseres som en ny form for viden i det kliniske arbejde. Med denne tilgang udviklede Iben i løbet af feltarbejdet en mere generel analytisk opmærksomhed på tid, og på hvordan kardiologerne etablerede ('enacted') forskellige tidsligheder i mødet med deres patienter, hvad enten de havde en prædiction tilgængelig eller ej. Den analytiske opmærksomhed på temporalitet bidrog til en forståelse af, hvordan algoritmens prognose spillede sammen med allerede eksisterende måder at praktisere fremtid på i behandlingen af patienter med iskæmisk hjertesygdom.

Artiklen her trækker primært på dette kliniske feltarbejde. Alle navne på informanter, forskningsprojektet Predict-IHD og CARDIA<sup>IHD</sup>-algoritmen er pseudonymer for at sikre anonymitet. Det etnografiske studie af algoritmen er godkendt af Datatilsynet. Dansk lov kræver ikke formel etisk godkendelse af sådanne studier.

## Analyse

I den følgende analyse viser vi gennem tre empiriske nedslagspunkter, hvordan CARDIA<sup>IHD</sup>-algoritmen – på trods af dens dokumenterede prædiktive performance – ikke nødvendigvis kunne informere og kvalificere kardiologernes beslutningstagning. Vores ærinde med at belyse og analysere disse tre empiriske udfordringer er ikke at kritisere eller underkende CARDIA<sup>IHD</sup>-algoritmen som konkret værktøj, men at undersøge og forstå AI-drevne prognoser som en ny vidensform, der søges integreret i etablerede kliniske praksisser.

### *Manglende data, fejlagtige prædictioner og lægers imputation af manglende datapunkter*

Danmark fremføres ofte som et af de mest digitaliserede og dataficerede lande i verden. I kraft af danske institutioner og myndigheders tradition for at indsamle data om hele befolkningen og danske borgere og patienters høje tillid til selvsamme myndigheder og institutioner, promoveres Danmark som et ideelt land for udvikling og implementering af nye dataintensive teknologier (Hoeyer, 2023). Håbet og ambitionen er, at prædiktive algoritmer trænet på komplette populationsdata af høj kvalitet kan sammenligne den enkelte patients livs- og sygdomsforløb med et stadigt voksende datagrundlag fra andre patienter og borgeres livsforløb – og i kraft af denne sammenligning forudsige med hidtil uset præcision, hvad fremtiden bringer for individet (Bauer, 2013). Antropologen Nick

Seaver argumenterer imidlertid for, at algoritmer mangler en sensitivitet over for den sociale, kulturelle og historiske kontekst, som værktøjerne opererer i (Seaver, 2015). Dermed kan der opstå en mulig kløft mellem en algoritmes prædiktions og de former for viden, mennesker anvender til at fortolke verden med, for eksempel intuition og viden om en patients sygdomshistorie (Anicini & Kotras, 2024). CARDIA<sup>IHD</sup>-algoritmens integration af historiske, patientspecifikke data søger at afhjælpe denne problematik ved at tage patientens helbreds- og sygdomsspecifikke kontekst i betragtning. I praksis viser det sig imidlertid sværere end som så at inddrage kontekst og forbinde individ og population i en præcis og meningsfuld forudsigelse.

En novemberdag i 2023 fulgte Iben kardiologen Hans, der arbejder på et højt specialiseret hospital i København. Den pågældende dag mødte hun en af Hans' patienter, Peter, en mand i midten af 50'erne fra Slagelse. Dagen før havde Peter fået en blodprop i hjertet. Umiddelbart inden ambulancen ankom til hospitalet i Slagelse, fik Peter hjertestop, og ambulancen blev derfor omdirigeret mod København og et mere specialiseret hospital. Undervejs i ambulancen fik Peter adskillige hjertestop og modtog både hjertemassage og elektriske stød på hjertet. Genoplivningsforsøgene lykkedes, og da Peter ankom til hospitalet i København, fik han lavet en ballonudvidelse, der genetablerede hjertets blodtilførsel. Før stuegang hos Peter orienterede Hans sig om hele forløbet i den elektroniske patientjournal. Her informerede et skærmvindue i Sundhedsplatformen om, at der var en CARDIA<sup>IHD</sup>-prædiktions tilgængelig for Peter. Med tanke på Peters blodprop og mange hjertestop i ambulancen var Iben overrasket over at se, at algoritmen klassificerede Peter som havende den bedste 1-års-overlevelse sammenlignet med gennemsnittet. Hans var til gengæld ikke overrasket. Han forklarede: "Algoritmen ved ikke, at han [Peter] har haft hjertestop, så derfor vrøvler den. Jeg vil estimere, at hans overlevelsesprognose er dårligere end gennemsnittet." I Peters tilfælde manglede algoritmen altså afgørende datainput fra ambulancen – enten fordi hjertestopdiagnosen ikke var blevet registreret korrekt, eller fordi algoritmen ikke kunne 'hente' dette datapunkt – og derfor beregnede den en alt for optimistisk overlevelsesprognose.

En kardiolog på et andet hospital i udkanten af København fortalte Iben en lignende historie om en patient fra Rumænien, der var blevet indlagt på deres afdeling med kardiogent chok og altså i en meget kritisk tilstand. Denne patient havde imidlertid ikke nogen 'historie' i det danske sundhedssystem, og derfor havde algoritmen kun meget få data at processere. Resultatet blev som i tilfældet med Peter, at algoritmen klassificerede den livstruende syge patient som havende den bedste overlevelsesprognose i forhold til gennemsnittet. Kardiologen sagde:

“Der skal ikke mange af disse fejl til, før man mister troen på algoritmen og lysten til at bruge den.” Når algoritmen af den ene eller anden årsag mangler kritiske datainput, oplever klinikerne en stærk kontrast mellem algoritmens prædiktioner og patienternes kliniske tilstand. I disse situationer fratager algoritmen ikke kardiologerne det komplekse arbejde med at integrere og fortolke de mange data og informationer, der er tilgængelige om en given patient.

I empiriske studier af AI og andre dataintensive teknologier har STS-forskere vist og begrebsliggjort, hvordan disse teknologier fører til nye former for ‘algoritmearbejde’ (Bailey et al., 2020), ‘dataarbejde’ (Bossen et al., 2019), ‘reparationsarbejde’ (Schwennesen, 2019) og ‘vedligeholdelsesarbejde’ (Lipp, 2022) for de menneskelige aktører. Tilsammen synliggør disse begreber, at nye teknologier og dataintensiv medicin ikke bare forenkler og overtager opgaver, men at de ofte giver anledning til, at opgaver og ansvar forflyttes, eller at nye opgaver opstår. De store forhåbninger om, at prædiktive algoritmer vil gøre det kliniske arbejde både nemmere og mere enkelt, har således vist sig at være vanskelige at indfri i praksis (se også Maiers, 2017; Rather & Elmholdt, 2023; Juestesen & Plesner, 2024). I tråd med disse fund viser vores observationer, at CARDIA<sup>IHD</sup>-algoritmen giver anledning til en ny form for fortolkningsarbejde, fordi kardiologen nu også skal vide, hvad algoritmen *ikke* ved, for at kunne vurdere rigtigheden og dermed den kliniske værdi af dens prædiktioner. I datavidenskaben betegner begrebet ‘imputation’, at manglende datapunkter udskiftes med en anden værdi, så den statistiske analyse kan inkludere alle individer i et datasæt og ikke kun dem med komplette datapunkter. Situationen med patienten Peter og patienten fra Rumænien uden data i sin elektroniske patientjournal eller danske registre viser, at det er kardiologerne fremfor algoritmen, der skal være i stand til at identificere og ‘imputere’ de manglende datapunkter – eller i hvert fald betydningen af deres fravær – som algoritmen af den ene eller anden grund ikke har til rådighed.

### *Fortolkning af korrelation som kausalitet*

Samtlige af de kardiologer, som Iben talte med under sit feltarbejde, udtrykte positive forventninger til at bruge AI i det kliniske arbejde og var entusiastiske i forhold til implementeringen af CARDIA<sup>IHD</sup>-algoritmen. De hilste mere data-drevne og dermed ‘objektive’ vurderinger af deres patienters fremtid velkomne. I praksis viste det sig imidlertid også at være udfordrende at mobilisere algoritmens prædiktioner i den kliniske beslutningstagning. Prædiktive algoritmer som CARDIA<sup>IHD</sup> finder statistiske sammenhænge i data, og det betyder, at de forklarende faktorer, som algoritmen lister op som drivende prædiktionen mod

enten overlevelse eller død, repræsenterer de variable, der statistisk er stærkest associeret med den beregnede prognose – ikke kausale forhold mellem de angivne faktorer og overlevelsesprognosen (cf. Amoore & Piotukh, 2016; Varga et al., 2020).

At algoritmen beregner korrelationer i data skaber udfordringer i den kliniske brug af den. Som eksempel fulgte Iben Hans i klinikken en efterårsdag i 2023, hvor han blandt andet gik stuegang hos en 72-årig mand, Anders, der var blevet indlagt med en blodprop i hjertet. Hans kunne læse i gamle journalnotater i Sundhedsplatformen, at Anders også havde diabetes, og at han var alvorligt syg af iskæmisk hjertesygdom: Han havde tidligere haft en blodprop i hjernen, han var før blevet opereret i hjertets kranspulsårer, og han havde fået amputeret det ene ben. Hans var derfor i dette tilfælde helt enig med CARDIA<sup>IHD</sup>-algoritmen, der kategoriserede Anders som havende den laveste 1-års-overlevelse i forhold til gennemsnittet. Da Hans kiggede nærmere på algoritmens forklarende faktorer, fremgik det, at algoritmen som forventet listede Anders' diagnoser relateret til åreforkalkning og diagnosen DE10 (Type 1-diabetes) som faktorer, der forringer overlevelsesprognosen. Hans anmærkede også, at diagnosekoderne DE36, (Forandringer i øjets nethinde ved sygdomme klassificeret andetsteds) og DE13, (Andre former for diabetes) drev prædiktionen mod overlevelse. Selvom Hans var helt klar over, at disse korrelationer gav mening statistisk – han kunne give kvalificerede gæt på hvordan men ikke forklare dem grundet modellens kompleksitet – så var de kontraintuitive fra et klinisk perspektiv. Iben observerede et lignende tilfælde hos en patient, hvor algoritmen angav normal koncentration af sukker i blodet, og at patienten ikke var kendt med skulderlidelser som faktorer, der drev prædiktionen mod død. Hans var betænkelig ved, at 'forklaringer' på modellens prædiktioner som disse kunne få nogle af hans kollegaer til ikke at bruge algoritmen, fordi man som kliniker kan være tilbøjelig til at afvise ikke-kausale sammenhænge, som man ikke forstår, eller som ikke passer med ens verdensbillede, og fordi klinikere efterspørger kausale sammenhænge, som de kan handle på.

I en tekstboks under algoritmens forklarende faktorer i Sundhedsplatformens visning fremgår det, at: "De forklarende faktorer kan ikke betragtes som kausale, men som korrelationer, som er identificerede i en stor patientgruppe med IHD [iskæmisk hjertesygdom]. Derfor kan der forekomme risikofaktorer, som virker ulogiske." Hans var dog ikke den eneste, der anmærkede, at algoritmens forklarende faktorer var svære at forstå. En septemberdag i 2024 var Iben på feltarbejde på et andet hospital i Region Hovedstaden. Under morgenkonferencen delte den kardiolog, der tidligere havde fortalt om sin oplevelse med den rumænske patient, sin erfaring med at bruge CARDIA<sup>IHD</sup>-

algoritmen: "De forklarende faktorer er for svære for os klinikere at forstå og bruge. Det kan godt være, at de giver mening statistisk, men de giver ikke mening klinisk." CARDIA<sup>IHD</sup>-algoritmen er et risikovurderingsværktøj, og formålet med prædiktionen er derfor ikke at give beslutningsstøtte til, hvilke risikofaktorer der bør påvirkes eller interverneres mod, men snarere at give et mere præcist billede af patientens samlede mortalitetsrisiko. På trods af, at kardiologerne bliver gjort opmærksomme på, at de forklarende faktorer netop er statistiske og ikke kausale sammenhænge, så kommer de, der ikke kender CARDIA<sup>IHD</sup>-algoritmen eller lignende værktøjer, nemt til at betragte de forklarende faktorer som nogle, der har en kausal forbindelse til overlevelsesprognosen. For klinikerne handler det om at identificere hvilke medicinske tiltag og interventioner, de kan igangsætte for at forbedre patientens situation. Her kan information og algoritmens parametre om patienten, som de ikke umiddelbart kan forstå eller handle på, for eksempel at en patient har forandringer i øjets nethinde, opleves som irrelevante eller som 'støj'. I den praktiske brug af algoritmen er kardiologerne tilbøjelige til at fortolke korrelation som kausalitet, og dette forhold gør algoritmen udfordrende at bruge for nogle kardiologer. For det første er det svært for klinikerne at forstå og acceptere korrelationer, der umiddelbart virker ulogiske fra et klinisk perspektiv. For det andet er det udfordrende for klinikerne ikke at fortolke de forklarende faktorer som nogle, de bør påvirke for at forbedre patientens tilstand. Algoritmer som CARDIA<sup>IHD</sup> aktualiserer dermed den velkendte problematik fra oversættelsen af epidemiologiske studier til klinisk praksis, hvor associationer let kan blive fortolket og handlet på som kausale forhold i kliniske vejledninger.

En diskussion fra den samme morgenkonference tjener som det sidste eksempel på, hvordan begrebet prognose for klinikerne er forbundet til deres muligheder for at intervenere medicinsk. En af de patienter, lægekollektivet diskuterede denne morgen, var en 90-årig kvindelig patient, der var blevet indlagt på afdelingen med et komplekst sygdomsbillede. Lægerne blev meget hurtigt enige om, at skulle kvinden blive rask, så ville det kræve en aggressiv og omkostningsfuld behandling, som hun ikke ville leve længe nok til at få gevinsten af – en situation som lægerne kalder 'det prognostiske paradoks'. Som afslutning på diskussionen og dermed beslutningen om ikke at igangsætte aggressiv behandling, konkluderede en læge: "Det er livskvalitet frem for prognose." For klinikerne er prognose altså tæt forbundet til deres muligheder for at igangsætte medicinske tiltag, der forlænger patientens liv på en måde, der er meningsfuld for både patienten og sundhedsvæsenet som helhed: Medicinske tiltag skal ikke kun resultere i et længere liv men også i livskvalitet i den tid, patienten har tilbage. Hvis patienten ikke har udsigt til et liv med livskvalitet, opfattes forlængelsen af liv – her

begrebsliggjort som prognose – som en byrde for patienten, der skal leve med mulige bivirkninger, og for sundhedsvæsenet og dermed samfundet, der spilder unødige og knappe ressourcer (Svendsen et al., 2018; Navne & Svendsen, 2017). I denne kliniske kontekst forstås og praktiseres prognose altså som en udsigt eller en levetid for patienten, som klinikeren skal kunne forandre og forlænge på en meningsfuld måde gennem konkret medicinsk intervention. Her kan CARDIA<sup>IHD</sup>-algoritmens kategoriske overlevelseshorisonter og de forklarende faktorer være svære for kardiologerne at mobilisere som retningsgivende viden for den kliniske beslutningstagning.

### *Multiple risici og fremtidshorisonter i klinisk praksis*

Under feltarbejdet erfarede Iben, at CARDIA<sup>IHD</sup>-algoritmens prognose blot var én ud af mange forskellige udsagn om fremtiden i den kliniske praksis. En efterårsdag i 2023 fulgte hun Hans, der som nævnt ovenfor er kardiolog på et højt specialiseret hospital, hvor kun de mest kritisk syge og komplicerede patienter bliver indlagt. Hans gik denne dag blandt andet stuegang hos en 65-årig mand, Arne, der var blevet indlagt dagen forinden med tiltagende brystmerter. Arne ventede på at få lavet enten en ballonudvidelse eller en bypass-operation den kommende dag. Under formiddagens stuegang sagde Hans til Arne:

*“Men hvis du får ondt i brystet før operationen, så skal vi vide det med det samme, for i det tilfælde vil de lave indgrebet akut [...] Nu tager vi én dag ad gangen. Vi skal have dig levende igennem vores del af behandlingen. Det ser lovende ud, men det er en alvorlig tilstand [...] Hvis indgrebet forløber ukompliceret, så flytter vi dig kort efter til dit hjemsygehus.”*

I dette korte uddrag af samtalen mellem Hans og Arne ser vi, at der er flere fremtider med dertil hørende risici på spil i den kliniske situation. Den første tidlige horisont, som Hans etablerer for Arne, rækker ikke længere end et døgnstid. Hans skal i første omgang sikre, at Arnes helbred ikke forværres, eller at han dør af sin hjertesygdom inden for de næste 24 timer og dermed før det planlagte indgreb. Med kommentaren ‘nu tager vi en dag ad gangen’ understreger Hans, at han først og fremmest skal håndtere de risici, der eksisterer i de kritiske timer efter en blodprop og komplicerede kliniske procedurer såsom en bypass-operation. Den næste horisont, Hans etablerer for Arne, omfatter dagene umiddelbart efter den planlagte operation, hvor Arne både er i risiko for at få nye blodpropper og forstyrrelser i hjertets elektriske ledningssystem, der i værste fald kan medføre hjertestop. Med formuleringen ‘vi skal have dig levende igennem vores del af

behandlingen', 'enacter' (Mol, 2002) Hans Arnes allernærmeste fremtid som både uforudsigelig og foranderlig. Set i forhold til disse to 'nære' fremtider præsenterer CARDIA<sup>IHD</sup>-algoritmen klinikerer for endnu en tidshorisont, nemlig patientens risiko for ikke at være i live efter et år. Formålet med denne AI-beregnete prognose er netop at bidrage til klinikerens risikovurdering i komplekse beslutningssituationer ved at placere patienten i forhold til det historiske datasæt, algoritmen er blevet trænet på, og dermed sammenligne patienten med den observerede overlevelse for lignende patienter. I den akutte fase, hvor nye komplikationer eller risikable indgreb kan ændre patientens tilstand radikalt, har klinikere ofte fokus på at håndtere de nære fremtider og risici. I praksis betyder det, at CARDIA<sup>IHD</sup>-algoritmens overlevelsesprognose, der er statisk og spår om en mere fjerntliggende fremtid, nemmere udgrænses eller træder i baggrunden.

I bogen *...And a Time to Die: How American Hospitals Shape the End of Life* demonstrerer antropologen Sharon Kaufman (2006), at behandling og omsorg for terminalt syge patienter er organiseret i 'pathways', og dermed hvordan sundhedsprofessionelles praksisser og erfaringer af tid former afslutningen af livet. Kaufman argumenterer for, at hospitalsinstitutionens krav om økonomisk og klinisk effektiv fremdrift gør timing af beslutninger og procedurer samt evnen til at håndtere opgaver til tiden til et dominerende anliggende for sundhedspersonalet: "Time is the marker for things health professionals think should happen and for things that must get done, and it weighs heavily on everyone who works in a hospital" (Kaufman, 2006, s. 7). Der er væsentlige forskelle mellem Kaufmans empiriske felt, der centrerer sig om behandling og omsorg for terminalt syge patienter, og det felt, som CARDIA<sup>IHD</sup>-algoritmen anvendes i, hvor kardiologer forsøger at redde liv ved pludseligt indsættende livstruende sygdom. Ikke desto mindre ser vi med Kaufmans analytik, at håndteringen af kritisk syge hjertepatienter ikke kan adskilles fra de temporale logikker og praksisser, som karakteriserer hospitalet som institution, og den måde det kliniske arbejde er organiseret på. For Hans, der arbejder på et højt specialiseret hospital, handler det frem for alt om at sikre, at patienter som Arne overlever de første par dage efter en operation, så de kan flyttes til det lokale hospital og dermed gøre plads til nye kritisk syge patienter. Den kliniske håndtering af Arnes risiko for at være død om et år er med andre ord forankret i andre og mere lokale institutioner.

Vi ser altså her, hvordan prognose – bredt forstået som en forudsigelse eller et udsagn om fremtiden – etableres på forskellige måder i det sociale møde mellem læge og patient i en given institutionel ramme (Jain, 2007; Gjødsbøl & Svendsen, 2019; Christakis, 2001). En analytisk opmærksomhed på de mange forskellige fremtider, der etableres i klinikken, viser også, at prognose ikke er et fænomen,

der kan afgrænses til det enkelte individ. Fordi etableringen af patienters 'pathways' er formet af det organisatoriske imperativ om fremdrift og effektivitet i det kliniske arbejde, bliver den enkelte patients fremtid(er) tæt vævet sammen med den måde, tid erfares og gøres på i en given klinisk kontekst. Det betyder også, at prognoser altid er situeret i både tid og sted. Eksemplet med Arnes forløb viser, at AI-drevne værktøjer til etablering af algoritmiske, personlige prognoser skal være nøje tilpasset den organisatoriske sammenhæng, som klinikerne står i, og de temporale horisonter, som er rammesættende for det kliniske arbejde, hvis de skal informere de kliniske beslutninger.

## Diskussion

Kardiologien – herunder behandling af akut hjertesygdom – er et højstatus medicinsk speciale, hvor brugen af algoritmer til risikostratificering som tidligere nævnt er velkendt og udbredt. Man kunne derfor forvente, at en algoritme som CARDIA<sup>IHD</sup> ville være forholdsvis ligetil at integrere i de kliniske ræsonnementer. Når dette ikke viste sig at være tilfældet, indikerer vores empiriske studie, at indføring og brug af komplekse algoritmer til andre medicinske felter og problemstillinger kan være endnu vanskeligere. For det første viser vores case, at selvom de mange og komplette danske populationsdata fra den brede befolkning i princippet muliggør prædiktion af den individuelle patients fremtid, så er datainfrastrukturer som CARDIA<sup>IHD</sup>-algoritmen skrøbelige, hvis de ikke har adgang til de nødvendige data i realtid og til data af tilstrækkelig kvalitet. Kardiologernes erfaring med og oplevelse af algoritmens upræcise prædiktioner eksemplificerer den mere generelle udfordring med brugen af AI i sundhedsvæsenet: at algoritmer har karakter af 'black boxes'. Black box refererer her til, at brugerne af en algoritme ikke kan gennemskue eller forstå, hvordan input bliver til et givent output (Pasquale, 2015; Gusterson, 2019; Mittelstadt et al., 2016). Manglen på transparens kan gøre, at en algoritme øger klinikernes oplevelse af usikkerhed frem for at reducere den. Denne nye form for epistemisk usikkerhed kan som vist i artiklen her og i et interviewstudie af Tyskbo og Nygren (2024) føre til 'handlingsmæssig usikkerhed' ("actionable uncertainty") frem for at give klinikerne en retning for hvilken beslutning, der bør træffes.

For det andet viser vores analyse, at det kan være svært for kardiologerne at forstå CARDIA<sup>IHD</sup>-algoritmen og at mobilisere den i deres beslutningstagning, fordi det kræver et fortolkningsmæssigt arbejde, som læger uden tilstrækkeligt kendskab til algoritmen ikke har forudsætningerne for at udføre. 'Explainable

AI fremføres i forskning og politik som løsningen på black box-problematikken beskrevet ovenfor og sigter på at øge gennemsigtigheden og brugernes forståelse af, hvorfor en algoritme når frem til et givent output (Arietta et al., 2020; Mittelstadt et al., 2019). CARDIA<sup>IHD</sup>-algoritmens forklarende faktorer er netop et sådant forsøg på at øge klinikernes forståelse af algoritmens overlevelsesprognoser. Vores empiriske cases viser dog, at denne transparens har en pris. I oversættelsen af den matematiske model til den kliniske kontekst kan det ske, at modellens beregninger af statistiske sammenhænge i data bliver fortolket som kausale sammenhænge. Hvis forklaringer på statistiske sammenhænge i data ikke giver mening fra et klinisk perspektiv, så får denne information ikke nødvendigvis karakter af en 'forklaring', der kan bidrage til klinikerens forståelse og fortolkning af prædiktionen og dermed være retningsgivende for hvilke kliniske beslutninger, der bør træffes. Med andre ord garanterer visningen af forklarende faktorer ikke nødvendigvis en korrekt og brugbar fortolkning af algoritmens output. I vores tilfælde havde Hans et indgående kendskab til CARDIA<sup>IHD</sup>-algoritmen og derfor gode forudsætninger for at forstå dens prædiktioner og de forklarende faktorer, samt hvad disse kan og ikke kan bruges til. Kardiologer, der ikke har kendskab til eller været involverede i Predict-IHD-projektet, har imidlertid ikke disse forudsætninger for at forstå algoritmen og vurdere dens prædiktioner derefter. Vi foreslår derfor en nytænkning af Marc Berg og Els Goormans (1999) stadig gyldige og højaktuelle 'Law of medical information', der foreskriver: "The further information has to circulate (i.e. the more different contexts it has to be usable in), the more work is required to disentangle the information from the context of its production" (Berg & Goorman, 1999, s. 51). Omsat til brugen af komplekse AI-drevne algoritmer må gælde, at jo større afstanden er mellem brugere og udviklere af komplekse algoritmer, desto mere fortolkningsmæssigt arbejde kræver det at forstå og bruge disse værktøjer på en sikker og meningsfuld måde.

For det tredje demonstrerer vores analyse betydningen af det kliniske arbejdes organisering i tid og rum for den måde, patienternes fremtider bliver etableret, kommunikeret og handlet på. I bogen *Death Foretold: Prophecy and Prognosis in Medical Care* undersøger læge og sociolog Nicholas A. Christakis (2001), hvordan læger forstår begrebet prognose og deres erfaring af at prognosticere for deres patienter. Christakis viser, at læger er tilbageholdende med at forudsige deres patienters fremtid – særligt døden – og at etableringen af sådanne forudsigelser er forbundet med svære følelser og etiske dilemmaer, fordi lægerne opfatter selve forudsigelsen som en aktør, der har betydning for, hvordan patientens fremtid udspiller sig. Med andre ord er lægerne i Christakis' studie bange for at fremskynde eller være en medvirkende årsag til en patients død ved at forudsige

den. Han argumenterer for, at på trods af den store usikkerhed, som karakteriserer forudsigelser, så er læger moralsk forpligtede på at italesætte og kommunikere deres bedste skøn om patientens fremtid. Uden denne viden, argumenterer Christakis, kan patienten og dennes pårørende ikke tage de fornødne forholdsregler, få mest muligt ud af den sidste tid eller planlægge livets afslutning, så den opfylder patientens og familiens ønsker. Som hjælp til at prognosticere hilser Christakis derfor nye datadrevne værktøjer velkommen. Analysen af de forskellige risici og dertilhørende tidshorisonter, som kardiologen etablerede og skulle håndtere for patienten Arne, viser imidlertid, at uanset hvor præcise, AI-drevne prædiktioner kan blive, så skal de passe ind i gældende organisatoriske krav om effektivitet og fremdrift – og dermed om at flytte patienter i tid og rum – for at blive meningsfulde og brugbare for klinikerne i praksis.

## Konklusion

I denne artikel har vi undersøgt den kliniske brug af CARDIA<sup>IHD</sup>-algoritmen, der afprøves som risikovurderingsværktøj for kardiologer, der behandler patienter med akut iskæmisk hjertesygdom. På trods af at algoritmen udelukkende anvendes på forsøgsbasis, giver vores etnografiske undersøgelse indsigter i, hvordan klinikere reagerer på AI-drevne værktøjer i sundhedsvæsenet. Vores analyse viser, at prædiktive algoritmer ikke kun fratager klinikere opgaver med at integrere og fortolke data, men også introducerer nye former for usikkerhed og kræver, at klinikerne engagerer sig i nye former for fortolkningsmæssigt arbejde for at kunne bruge algoritmen efter hensigten og forstå dens begrænsninger. Derudover peger vores undersøgelse på, at klinikere kan have udfordringer med at afkode, hvad de kan bruge en algoritmes prædiktioner samt eventuelle forklarende faktorer til, samt hvilken adfærd der forventes af dem, når de præsenteres for algoritmens prædiktion. Disse praktiske udfordringer i den kliniske brug af CARDIA<sup>IHD</sup>-algoritmen vidner om, hvor komplekst – både teknisk, praktisk, organisatorisk og kulturelt – det er at udvikle og implementere nye AI-drevne værktøjer i klinikken. På den ene side kan algoritmer som CARDIA<sup>IHD</sup> betragtes som en intensivering af risikosamfundets logik (Beck, 1992), idet de producerer stadig mere individualiseret risikovidens frem for risikovidens på befolknings- og gruppeniveau. På den anden side viser vores etnografiske analyse, at de 'personlige prognoser', som genereres af dataintensive, algoritmiske værktøjer, kan skabe nye fortolkningsmæssige udfordringer og være vanskelige for klinikere at omsætte til konkrete beslutninger. Vores studie peger

dermed på, at udbredelsen af algoritmisk risikovidens kan overstige de praktiske handlemuligheder for at omsætte den i praksis.

Afslutningsvist er det vigtigt at understrege, at de udfordringer med at bruge CARDIA<sup>IHD</sup>-algoritmen i praksis, som artiklen har belyst, primært handler om, hvordan klinikere præsenteres for AI-drevne løsninger i praksis og ikke så meget om, hvorvidt algoritmen fungerer som matematisk model. Algoritmen er som nævnt videnskabeligt valideret som et mere præcist risikovurderingsværktøj i forhold til allerede eksisterende værktøjer, og der kan først konkluderes på dens kliniske effekt, når resultaterne af det kliniske lodtrækningsforsøg foreligger. Uanset hvordan dette studie falder ud, så repræsenterer CARDIA<sup>IHD</sup>-algoritmen et nybrud som datainfrastruktur, der gør det muligt at arbejde med data-integration, automatisering og risikovurderinger i klinisk praksis på en ny måde. CARDIA<sup>IHD</sup>-algoritmen kan dermed være med til at bane vejen for udvikling og implementering af nye AI-drevne løsninger i det danske sundhedsvæsen.

## Referencer

- Akrich, M. (1992). The De-Description of Technical Objects. In W.E. Bijker & J. Law (Eds.), *Shaping Technology/Building Society: Studies in Sociotechnical Change*. (pp. 205-225). Cambridge, MA: MIT Press.
- Ali, O., Abdelbaki, W., Shrestha, A., Elbasi, E., Alryalat, M.A.A., & Dwivedi, Y.K. (2023). A systematic literature review of artificial intelligence in the healthcare sector: Benefits, challenges, methodologies, and functionalities. *Journal of Innovation & Knowledge*, 8(1), 100333. <https://doi.org/10.1016/j.jik.2023.100333>.
- Amelang, K., & Bauer, S. (2019). Following the Algorithm: How Epidemiological Risk-Scores Do Accountability. *Social Studies of Science*, 49(4), 476-502. <https://doi.org/10.1177/0306312719862049>.
- Amoore, L. (2013). *The politics of possibility. Risk and security beyond probability*. Durham: Duke University Press.
- Amoore, L., & Piotukh, V. (2016). Introduction. In L. Amoore & V. Piotukh (Eds.), *Algorithmic Life: Calculative devices in the age of big data*. (pp. 1-18). Abingdom and New York: Routledge.
- Anichini, G., & Kotras, K. (2024). Trained Judgements: Artificial Intelligence, Epistemic Tensions and the Production of Scientific Objectivity. *Science, Technology, & Human Values*, 51(3), 631-663. <https://doi.org/10.1177/01622439241262854>.
- Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., et al. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115. <https://doi.org/10.1016/j.inffus.2019.12.012>.
- Backus, B.E., Six, A.J., Kelder, J.C., et al. (2010). Chest pain in the emergency room: a multicenter validation of the HEART score. *Critical Pathways in Cardiology*, 15(2), 40-15. <https://doi.org/10.1097/HCP.0000000000000070>.

- Bailey, S., Pierides, D., Brisley, A., Weisshaar, C., & Blakeman, T. (2020). Dismembering organisation: The coordination of algorithmic work in healthcare. *Current Sociology*, 68(4), 546-571. <https://doi.org/10.1177/0011392120907638>.
- Bajwa, J., Munir, U., Nori, A., & Williams, B. (2021). Artificial intelligence in healthcare: transforming the practice of medicine. *Future Healthcare Journal*, 8(2), 188-194. <https://doi.org/10.7861/fhj.2021-0095>.
- Bauer, S. (2013). Modeling Population Health. *Medical Anthropology Quarterly*, 27, 510-530. <https://doi.org/10.1111/maq.12054>.
- Beck, U. (1992). *Risk society: Towards a new modernity*. London: Sage
- Berg, M. (1997). *Rationalizing medical work: Decision support techniques and medical practice*. Cambridge, MA: MIT Press.
- Berg, M., & Goorman, E. (1999). The contextual nature of medical information. *International Journal of Medical Information*, 56(1-3), 51-60. [https://doi.org/10.1016/s1386-5056\(99\)00041-6](https://doi.org/10.1016/s1386-5056(99)00041-6).
- Bossen, C., Pine, K.H., Cabitza, F., Ellingsen, G., & Piras, E.M. (2019). Data work in healthcare: An Introduction. *Health Informatics Journal*, 25(3), 465-474. <http://doi.org/10.1177/1460458219864730>.
- Christakis, N. (2001). *Death Foretold: Prophecy and Prognosis in Medical Care*. Chicago, IL: Chicago University Press.
- Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal*, 6, 94-98. <http://doi.org/10.7861/futurehosp.6-2-94>.
- D'Ascenzo, F., Biondi-Zoccai, G., Moretti, C., et al. (2012). TIMI, GRACE and alternative risk scores in Acute Coronary Syndromes: a meta-analysis of 40 derivation studies on 216,552 patients and of 42 validation studies on 31,625 patients. *Contemporary Clinical Trials*, 33(3), 507-514. <http://doi.org/10.1016/j.cct.2012.01.001>.
- Dixon, D., Sattar, H., Moros, N., Kesireddy, S.R., Ahsan, H., Lakkimsetti, M., Fatima, M., Doshi, D., Sadhu, K., & Junaid Hassan, M. (2024). Unveiling the Influence of AI Predictive Analytics on Patient Outcomes: A Comprehensive Narrative Review. *Cureus*, 16(5), e59954. <http://doi.org/10.7759/cureus.59954>.
- Eftekhari, A., Sadjadieh, G.C., Kyhl, K., Pareek, M., Jacobsen, M.R., Hougaard, M.V., Støttrup, N.B., Jabbari, R. & Fakhri, Y. (2025). Akut koronart syndrom. Dansk Cardiologisk Selskab. Tilgængelig på: *Behandlingsvejledning | Akut koronart syndrom*. [Tilgæet den 27. januar 2026].
- Esposito, E. (2024). Can we use the open future? Preparedness and innovation in times of self-generated uncertainty. *European Journal of Social Theory*, 27(2), 209-224. <https://doi.org/10.1177/13684310231224546>.
- Esposito, E., Hofmann, D., & Coloni, C. (2024). CAN A PREDICTED FUTURE STILL BE AN OPEN FUTURE? ALGORITHMIC FORECASTS AND ACTIONABILITY IN PRECISION MEDICINE. *History and Theory*, 63, 4-24. <https://doi.org/10.1111/hith.12327>.
- Fabian, J. (1983). *Time and the Other: How Anthropology Makes Its Object*. New York: Columbia University Press.
- Frank, L. (2024). Mit personlige orakel. Tilgængelig på: *Mit personlige orakel | Ideer | Weekendavisen*. [Tilgæet den 8. april 2026].
- Gjødsvøl, I.M. & Svendsen, M.N. (2019) Time and Personhood across Early and Late-Stage Dementia. *Medical Anthropology*, 38(1), 44-58. 10.1080/01459740.2018.

- Gusterson, H. (2019). Introduction: Robohumans. In Besteman, C., & Gusterson, H. (Eds.), *Life by Algorithms: How Roboprocesses Are Remaking Our World*. (pp. 1-27). Chicago and London: The University of Chicago Press.
- Haug, C.J., & Drazen, J.M. (2023). Artificial Intelligence and Machine Learning in Clinical Medicine. *The New England Journal of Medicine*, 388(13), 1201-1208. <https://doi.org/10.1056/NEJMra2302038>.
- Hoeyer, K. 2023. *Data Paradoxes: The Politics of Intensified Data Sourcing in Contemporary Healthcare*. Cambridge, MA: MIT Press.
- Jain, S.L. (2007) Living in prognosis: Toward an Elegiac Politics. *Representations*, 98(1), 77–92. <https://doi.org/10.1525/rep.2007.98.1.77>.
- Johnson, K.B., Wei, W.Q., Weeraratne, D., Frisse, M.E., Misulis, K., Rhee, K., Zhao, J., & Snowdon, J.L. (2021). Precision Medicine, AI, and the Future of Personalized Health Care. *Clinical and Translational Science*, 14(1), 86-93. <http://doi.org/10.1111/cts.12884>.
- Justesen, L., & Plesner, U. (2024). Angry Citizens and Black Belt Employees: Cascading Classifications of and around a Predictive Algorithm. *Valuation Studies*, 11(1), 9-37. <https://doi.org/10.3384/V5.2001-5992.2024.11.1.9-37>.
- Kaufman, S.R. (2006). *...And a Time to Die: How American Hospitals Shape the End of Life*. Chicago, IL: Chicago University Press.
- Koselleck, R. (1985). *Futures Past: On the Semantics of Historical Time*. New York: Columbia University Press.
- Knuuti, J., Wijns, W., Saraste, A., et al. (2020) ESC Guidelines for the diagnosis and management of chronic coronary syndromes. *European Heart Journal*, 41(3), 407-477. <https://doi.org/10.1093/eurheartj/ehz425>.
- Lebovitz, S., H. Lifshitz-Assaf, & Levina, N. (2022). To Engage or Not to Engage with AI for Critical Judgments: How Professionals Deal with Opacity When Using AI for Medical Diagnosis. *Organization Science*, 33(1), 126-148. <https://doi-org/10.1287/orsc.2021.1549>.
- Lipp, B. (2022). Caring for robots: How care comes to matter in human-machine interfacing. *Social Studies of Science*, 53(5), 660-685. <https://doi.org/10.1177/03063127221081446>.
- Maiers, C. (2017). Analytics in Action: Users and Predictive Data in the Neonatal Intensive Care Unit. *Information, Communication & Society*, 20(6), 915-929. <https://doi.org/10.1080/1369118X.2017.1291701>.
- Mattingly, C. (1998). *Healing Dramas and Clinical Plots: The Narrative Structure of Experience*. Cambridge, UK: Cambridge University Press.
- Mittelstadt, B.D., Russell, C., & Watcher, S. (2019). Explaining explanations in AI. *Proceedings of the 2019 conference on fairness, accountability, and transparency*, Atlanta, GA, United States, 279-288. <https://doi.org/10.1145/3287560.3287574>.
- Mol, A. (2002). *The Body Multiple: Ontology in Medical Practice*. Durham, NC: Duke University Press.
- Navne, L.E., & Svendsen, M.N. (2017). Careography: Staff Experiences of Navigating Decisions in Neonatology in Denmark. *Medical Anthropology*, 37(3), 253-266. <https://doi.org/10.1080/01459740.2017.1313841>.
- Obermeyer, Z., & Lee, T. (2017). Lost in thought: The limits of the human mind and the future of medicine. *New England Journal of Medicine*, 377(13), 1209-1211. <https://doi.org/10.1056/NEJMp1705348>.
- Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, NY: Harvard University Press.
- Prainsack, B. (2017). *Personalized Medicine: Empowered Patients in the 21st Century?* New York: New York University Press.

- Ratner, H.F., & Elmholdt, K. (2023). Algorithmic constructions of risk: Anticipating uncertain futures in child protection services. *Big Data & Society*, 10(2), 1-12. <https://doi.org/10.1177/20539517231186120>.
- Robusthedskommissionen. (2023). Robusthedskommissionens anbefalinger. Tilgængelig på: *Robusthedskommissionens anbefalinger – September 2023*. [Hentet den 3. juni 2026].
- Schwennesen, N. (2019). Algorithmic Assemblages of Care: Imaginaries, Epistemologies and Repair Work. *Sociology of Health and Illness*, 41, 176-192. <https://doi.org/10.1111/1467-9566.12900>.
- Schwennesen, N., Svendsen, M.N., & Koch, L. (2011). Beyond informed choice: prenatal risk assessment, decision-making and trust. *Clinical Ethics*, 5(4), 207-216. 10.1258/ce.2010.010041.
- SCORE2 working group & ESC Cardiovascular risk collaboration (2021). SCORE2 risk prediction algorithms: new models to estimate 10-year risk of cardiovascular disease in Europe. *European Heart Journal*, 42(25), 2439-2454. <https://doi.org/10.1093/eurheartj/ehab309>.
- Seaver, N. (2015). The Nice Thing About Context Is That Everyone Has It. *Media, Culture & Society*, 37(7), 1101-1109. <https://doi.org/10.1177/0163443715594102>.
- Sundhedsstrukturkommissionen. (2024). Sundhedsstrukturkommissionens rapport: Beslutningsgrundlag for et mere lige, sammenhængende og bæredygtigt sundhedsvæsen. Tilgængelig på: *Sundhedsstrukturkommissionens rapport*. [Hentet den 3. juni 2026].
- Svendsen, M.N., Gjødsbøl, I.M., Navne, L.E., & Dam, M.S. (2018). A Life Worth Living: Temporality, Care, and Personhood in the Danish Welfare State. *American Ethnologist*, 45(1), 20-33. <https://doi.org/10.1111/amet.12596>.
- Tavory, I., & Timmermans, S. (2014). *Abductive Analysis: Theorizing Qualitative Research*. Chicago: The University of Chicago Press.
- Timmermans, S. (1996). Saving Lives or Saving Multiple Identities?: The Double Dynamic of Resuscitation Scripts. *Social Studies of Science*, 26(4), 767-797. doi: 10.1177/030631296026004003.
- Topol, E. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25, 44-56. <https://doi.org/10.1038/s41591-018-0300-7>.
- Tyskbo D., & Nygren, J. (2024). Reconfiguration of uncertainty: Introducing AI for prediction of mortality at the emergency department. *Social Science & Medicine*, 359, 117298. <https://doi.org/10.1016/j.socscimed.2024.117298>.
- Varga, T.V., Niss, K., Estampador, A.C., Collin, C.B., & Moseley, P.L. (2020). Association is not prediction: A landscape of confused reporting in diabetes – A systematic review. *Diabetes Research and Clinical Practice*, 170, 108497. <https://doi.org/10.1016/j.diabres.2020.108497>.
- Willerslev, R., Christensen, D.R., & Meinert, L. (2013). *Introduction. I: Taming Time, Timing Death: Social Technologies and Ritual*. D. R. Christensen & R. Willerslev, eds. (Pp. 1–16). Farnham, UK: Ashgate.
- Yu, K.-H., Beam, A.L., & Kohane, I.S. (2018). Artificial intelligence in healthcare. *Nature Biomedical Engineering*, 2, 719-731. <https://doi.org/10.1038/s41551-018-0305-z>.

# Collective efforts for the introduction of AI in healthcare

## An empirical study of two Norwegian initiatives

Mari Serine Kannelønning<sup>1</sup>  
Miria Grisot<sup>1</sup>

---

<sup>1</sup>Institutt for Informatikk, Universitet i Oslo  
[mkanneloenning@gmail.com](mailto:mkanneloenning@gmail.com)  
[miriag@ifi.uio.no](mailto:miriag@ifi.uio.no)

Kannelønning, Mari Serine & Grisot, Miria. 2026. Collective efforts for the introduction of AI in healthcare: An empirical study of two Norwegian initiatives. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44, 100-123. DOI: 10.7146/tfss.v25i44.156499

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

*The collective dimension of AI is crucial to ensure sustainable and ethical outcomes of AI implementation in society; accordingly, the introduction of AI in healthcare is a collective concern. In this paper, we examine this collective dimension of AI introduction by drawing on an empirical longitudinal case study of two initiatives aimed at introducing AI into the Norwegian public healthcare services. The first initiative is an inquiry process led by the Norwegian Directorate of Health, as part of a project initiated by a government policy plan outlining a vision for AI in healthcare. The second initiative is a nationwide, informally established network of professionals engaged with AI in healthcare. Drawing on insights on hybrid knowledge production in collective efforts, as well as practice-oriented and relational understandings of technology, our analysis of the two initiatives focuses on how the collective effort was carried out, the role of the orchestrators and the implications*

*of such efforts on their outcomes. Our findings show that the two initiatives, while seeking to produce hybrid knowledge on how to best introduce AI in healthcare, faced different challenges, both in generating such knowledge, in achieving consensus-based results and in enacting visions of AI in healthcare oriented to clinical practice. We suggest that to contribute to a consensus-based and inclusive introduction of AI in healthcare, collective efforts should pay attention to the inclusion of heterogeneous stakeholders and their interaction modes, the politics of agenda settings and modes of orchestration.*

## Introduction

The introduction of AI in healthcare refers to the implementation, adoption and use of AI-based technologies in healthcare practices, and it is currently the focus of much scholarly and public discourse. Much of this discourse is centred around technological efficiency and performance, and it is preoccupied with understanding how advancements in AI will enhance the diagnostic accuracy and operational efficiency of healthcare professionals (see, for example, national policies like the Danish ‘National Strategy for Artificial Intelligence’ [Danish Ministry of Finance & Danish Ministry of Industry, 2019] or the Norwegian ‘Nasjonal helse- og samhandlingsplan 2024–2027’ [Meld. St. 9 (2023–2024)]). However, AI technologies are not a set of discrete technology devices being implemented to improve isolated healthcare tasks; rather, they have the potential to profoundly transform healthcare practices, services and systems (Alami et al., 2020; Noorbakhsh-Sabet et al., 2019; Periañez et al., 2024; Alowais et al., 2023). It is then critical to examine how this transformation unfolds, especially in the public healthcare sector of the Nordic countries.

In the Nordics, policy-making and societal development are traditionally characterised by consensus-building, dialogue, inclusion, and collective action (Kuhnle & Alestalo, 2017, p. 18; Lægreid, 2017). These are reflected in the current approach to the introduction of AI. In Norway, for instance, the government’s vision of AI, as presented in the Nasjonal helse- og sykehusplan 2020–2023 (Meld. St. 7 (2019–2020)), emphasises that enabling AI in healthcare requires a broad mobilisation of actors and public-private collaboration (p. 89). Although processes striving for consensus are typically time-consuming, the results are often less radical yet more sustainable than those of faster-moving reformers, such as countries like the UK (Lægreid, 2017, p. 85). Moreover, scholars within AI ethics underscore that collective approaches are of great importance in achieving “a good AI society” (Floridi et al., 2018). Floridi et al. argue that a multi-stakeholder

approach “is the most effective way to ensure that AI will serve the needs of society, by enabling developers, users and rule-makers to be on board and collaborating from the outset” (Floridi et al., 2018, p. 701). However, in an AI in healthcare context, knowledge of how such collective work can be done to ensure sustainable, ethical, and otherwise ideal societal outcomes is still limited. In this paper, we aim to explore how a collective approach to AI can be enacted. Our research is guided by the following research question: *How can collective efforts be organised and carried out to contribute to a consensus-based and inclusive introduction of AI in healthcare?*

To address this research question, we have conducted a longitudinal case study focusing particularly on two initiatives characterised by a collective approach in the early phases of AI introduction in Norway. The first initiative is an inquiry process led by the Norwegian Directorate of Health as part of a project initiated by a government policy plan outlining a vision for AI in healthcare published in 2019 (Meld. St. 7 (2019–2020)). The second initiative is an open nationwide network of professionals aiming to enable AI in healthcare. The initiatives were followed for five months and two years, respectively, using an ethnographic approach and data collection methods including meeting observations and interviews. Our findings indicate that the two initiatives encountered distinct challenges stemming from their differing collective approaches. The study contributes to the understanding of how the collective dimension of AI can play out, and our findings have implications for future initiatives that aim to collectively produce knowledge to better comprehend the challenges associated with emerging technologies and societal change.

The paper is structured as follows: It starts by introducing relevant literature on hybrid knowledge production and collective efforts, followed by a description of the two cases studied. The methods and materials are then accounted for, whereafter the findings are presented and subsequently discussed. Lastly, a section on the final remarks is included.

## Hybrid knowledge production and collective efforts

This paper focuses on collective efforts for introducing AI in healthcare. By collective efforts, we mean initiatives in which multiple diverse actors are organised and coordinated, share a certain degree of interest, and exercise collective agency. In collective action, actors coordinate their activities towards an outcome that would not be possible for any of them to achieve alone.

The dynamics of collective efforts have received much attention in research. Classic studies have, for instance, addressed the fundamental and complex tension between individual and collective interests. The seminal work of Olson (1965) on the logic of collective action questions the ability of groups to realise individuals' interests. This work points to the role of cooperation and coordination among the members of the collective, arguing that smaller groups of homogeneous actors are more likely to form collectives and achieve their goals (Olson, 1965). The underlying argument is that in large groups, actors find it more challenging to see their own benefits realised, and therefore they lose interest in contributing to the group. The concept of 'tragedy of the commons', for instance, indicates a dynamic in which individuals are driven by their own benefits in using common resources, such as natural resources, and do not consider the cumulative impact of their actions (Hardin, 1968). Thus, individuals' interests are put above the group's shared interests, resulting in unfavourable outcomes. Research has proposed ways to address this problem. Ostrom studied collective action and the role of institutions in the governance of common goods (Ostrom, 1990). Her research shows that self-organising collectives can govern common resources without centralised regulations and strong institutions and overcome the "tragedy". Thus, she argues for the importance of social interaction to address social dilemmas and shows that heterogeneous groups can achieve their goals by focusing on long-term solutions (Ostrom, 2005).

Other research has investigated how knowledge is produced in collective efforts. According to Gibbons et al. (1994), when a heterogeneous group of actors is involved, knowledge is produced through negotiations and interactions in which the actors' heterogeneity is a constitutive element. Negotiations allow for reflexivity and higher transparency as actors question each other's perspectives, underlying values, and visions (Gibbons et al., 1994). According to Callon, Lascoumes, and Barthe (2011), collectives can be viewed as hybrid forums that blend various forms of interaction among heterogeneous actors, including laypeople and experts, acts and values, and humans and non-humans, representing diverse interests and communities. They are also hybrid because "the questions and problems taken up are addressed at different levels in a variety of domains" (Callon et al., 2011, p. 18). In doing so, hybrid forums facilitate knowledge sharing and the production of hybrid knowledge (Gibbons et al., 1994) which is considered necessary for understanding and beginning to address the uncertainties typically associated with technological advancements (Callon et al., 2011).

Collective efforts to produce hybrid knowledge are fundamentally shaped by, but also potentially constrained by, the very constellations of actors that

constitute them. These actor constellations are central in relational and practice-based theories within Science and Technology Studies (STS), which hold that what technologies “are” depends on how they are enacted by various actors in concrete practices. For instance, Annemarie Mol (2002) shows that there are not only multiple perspectives on a given phenomenon or object within certain group constellations. Most importantly, there is ontological multiplicity: Actors enact technologies differently in practice based on their realities. In addition, technologies are envisioned and made to work according to the resources available in a specific contextual practice, where different understandings of a working technology emerge. Thus, what is considered technology success or failure varies (de Laet & Mol, 2000). Furthermore, in this view, as humans and machines are not fixed entities, and as what they are and their roles are continually produced and negotiated in practice, this ongoing interaction also produces and changes the practices (Suchman, 2007).

Additionally, the outcome of collective efforts is shaped by their orchestrators (Reypens et al., 2021). Reypens et al. (2021, p. 62) conceptualise three modes of orchestration: consensus-based, dominating and hybrid (the latter being a combination of the former two). They argue that while consensus-based orchestration works in small collectives, it can be challenging for the orchestrators to discern how relationship-building and knowledge production unfold when the number of stakeholders involved is high. Larger collectives can become opaque, making good communication and proper collective work harder. In such a case, a dominating orchestration might be a better option (Reypens et al., 2021). However, if the network consists of stakeholders with a wide range of expertise, dominating orchestrators will not have the same level or breadth of expertise and will face difficulties in maintaining legitimacy. Thus, Reypens et al. suggest that hybrid orchestration may be a solution for managing both many stakeholders and their diversity to achieve inclusive collective efforts and outcomes (2021, p. 62). By hybrid orchestration, Reypens et al. refer to orchestrators switching between a dominating mode and a consensus-based mode during different phases. For instance, during phases of connecting stakeholders in a network and initiating the first meetings, a dominating orchestration could be beneficial. Differently, when aiming to facilitate bottom-up collaboration and relationship-building, a consensus-based orchestration could be more appropriate. Moreover, when effective outcomes and concrete results are in focus, a dominating orchestration mode could again be more successful (Reypens et al., 2021, pp. 69, 76).

Based on this conceptual ground, we analyse two initiatives aimed at contributing to the introduction of AI in the Norwegian public healthcare services.

In doing so, the paper aims to advance our understanding of AI introductions as collective, negotiated and situated processes.

## Case description

This paper is based on a study of two key collective initiatives in the early phases of introducing AI into Norwegian public healthcare services, prior to the implementation of notable AI deployments in real-world clinical settings.

### *The national inquiry process (case 1)*

The first initiative, *the national inquiry process*, was carried out in 2021 and led by the Norwegian Directorate of Health. It aimed to gather insights from multiple stakeholders on how to enable AI in healthcare.

In November 2019, the first Norwegian policy plan addressing AI in public healthcare services was published (the Norwegian National Health and Hospital Plan (NHSP) 2020-2023). It contained the first formally articulated governmental vision of a future with AI in the healthcare services:

*“Artificial intelligence makes it possible to utilise national medical databases to provide faster and more accurate diagnostics, better treatment and a more effective use of resources. The regulations make it possible to use medical data to bring maximum benefit to the community, the directorates provide guidance on legal restrictions, and ethical problems associated with the use of artificial intelligence are handled in cooperation with other European countries. The health and care service has established a culture of innovation and knowledge-sharing, and works closely with businesses to develop the tools needed by the service and patients” (Meld. St. 7 (2019–2020), p. 26).*

The policy plan also stated that relevant government agencies should coordinate efforts to enable AI in healthcare (Meld. St. 7 (2019–2020), p. 98). Consequently, a national cross-public-agency coordination project was established. The agencies involved included the Directorate of eHealth, the Norwegian Medicines Agency, the four Norwegian regional health authorities and the Directorate of Health, which was set to manage the project. As one of its first activities, the project team initiated an inquiry process (case 1) to identify challenges and recommend actions to ensure the safe and appropriate implementation and use of AI in the healthcare services (Helsedirektoratet et al., 2021, p. 2).

The inquiry process lasted for five months (January-May 2021) and involved approximately 100 stakeholders with diverse backgrounds and expertise, including AI researchers, vendors, project managers, physicians, hospital managers, procurement officers and patient organisations. By focusing on radiology – a medical area that was (and still is) at the forefront of AI development and implementation – the project team expected to generate knowledge and experience that would be relevant and transferable to other medical areas. Three types of meetings were carried out during the inquiry process: insight meetings, workshops and feedback meetings. As the process took place during the Covid-19 pandemic, the meetings were all held digitally. The number of participants at each meeting typically ranged from 10 to 25, including approximately three to four members of the project team (details about the case are also reported in Kannelønning, 2023).

### *The nationwide network (case 2)*

The second initiative, *the nationwide network*, was established in 2020 and led by a secretariat. It aimed to bridge what they called the ‘implementation gap’ in clinical AI – the gap between existing technologies and their limited implementation – by facilitating knowledge and experience sharing among its members and influencing relevant decision-makers.

The network was initiated by a group of eight individuals with shared interests in AI in healthcare, including a private company researcher, a public agency department manager, two research centre managers, and a manager and a senior adviser from two different regional health authorities. They formed the secretariat that organised the network and facilitated its activities in the years that followed. According to a member of the secretariat, the network was founded on the belief that enabling AI in healthcare requires collaboration among people with diverse expertise and worldviews, and that introducing AI entails addressing technical, organisational, legal, ethical, and financial issues. The network aimed to accelerate the introduction of AI by bringing together professionals with relevant expertise to share knowledge and experience. Furthermore, the network aimed to influence relevant decision-makers based on the knowledge it produced. This informal ‘bottom-up’ collective effort and its goal of affecting the agenda of crucial decision-makers were emphasised as fundamental and explained by an informant as follows:

*“It’s quite simple: Keeping the focus on what the people with knowledge of AI in healthcare perceive as difficult or useful, what the solutions to the problems are, and*

*so on, will result in a more professional-oriented agenda than a political one. If we [the network] find different potential [in AI], which the top management either doesn't see, doesn't include in strategies, or doesn't make decisions about, then you have to do what you normally do: Try to let it [the knowledge] trickle upwards in the system, talk to your bosses and so on"* (a secretariat member, our translation).

The main activities of the network were quarterly digital meetings and three or four in-person seminars per year. The first official meeting took place in early 2021 and had 20 participants from across the country. With a membership policy stating that it was open to anyone interested, the network steadily grew over two years to include approximately 200 professionals from different fields, organisations and sectors with an interest in AI in healthcare. The largest group of members was researchers from the hard sciences (e.g., informatics and medical physics) from universities, university hospitals, or research centres. The second-largest group consisted of hospital employees (e.g., from IT, research, or procurement departments), followed by bureaucrats with backgrounds in medicine and/or technology who worked on digitalisation and innovation in healthcare. The remaining members were representatives from software companies, interest organisations and municipalities. Despite these professional roles, the membership was emphasised as personal, meaning members were expected to express their personal opinions rather than represent the official standpoints of their affiliated organisation. Additionally, as of February 2023, the network had 32 listed observers. This included the first author of this paper<sup>1</sup>, project managers, and senior advisers affiliated with organisations such as the Norwegian Directorate of Health, the Norwegian Board of Health Supervision and the Norwegian Board of Technology (details about the case are also reported in Kannelønning, 2024a)

## Methodology, data collection, and analysis

The study followed a longitudinal qualitative case study design. The first author collected data through observation of digital meetings, document analysis and online interviews as fieldwork was conducted during the pandemic. More

---

1. After emailing a member of the secretariat, the first author obtained permission to observe the network for research purposes. She is a social scientist affiliated with a Norwegian university and does not work in the healthcare sector. The research project and data management practice, including anonymisation procedures and data storage, was approved by SIKT – the Norwegian Agency for Shared Services in Education and Research. The interview informants were informed about the procedures and signed a consent form.

specifically, knowledge of the inquiry process (case 1) was gained through non-participant observations of eight stakeholder meetings (of approx. two hours each), content analysis of five reports produced by the inquiry process project team, and three meetings with the project manager. The knowledge of the nationwide network (case 2), its organisation, meetings, and activities were similarly gained through non-participant meeting observations of ten network meetings (of approx. three hours) during a two-year period (2021-2022) and analysis of documents like PowerPoint presentations from the network meetings and the network's website. Six of the secretariat members were also interviewed (see table 1 for details on the data collection, methods, and materials). While these interviews provided knowledge about the network as a collective effort from the organisers' perspective, the meeting observations of both initiatives provided insights into how the collective effort unfolded among the stakeholders involved.

*Table 1: Data collection: meetings, documents, interviews*

	<b>Case 1</b>	<b>Case 2</b>	<b>Total</b>
Number of meetings and approx. duration	8 meetings 17 h 30 m	10 meetings (incl. conferences) 57 h	18 meetings 84 h 30 m
Documents analysed	5 documents 152 p.	-	5 documents 152 p.
Interviews and approx. duration	-	6 interviews 6 h	6 interviews 6 h
Meetings with project manager	3 meetings approx. 2 h		3 meetings approx. 2 h

The data were analysed using an abductive approach, moving back and forth between data and theoretical concepts (Timmermans & Tavory, 2012). We started the process by immersing ourselves in the data, reading the material line by line and identifying and highlighting content relevant to the theme of collective efforts. As we extracted data from fieldnotes, interview transcripts, and documents, we began writing empirical narratives while engaging with the core concepts of hybrid knowledge, hybrid forums, and orchestrator modes. We iteratively revisited the empirical narratives, guided by the concepts, and re-read the relevant literature to interpret the concepts guided by our data. This process

helped us identify additional details in the material and analytical points that were not apparent during the initial reading phase. For instance, guided by the literature on orchestration modes, we became aware of the organisers' role in facilitating knowledge production within collective efforts in different ways and across phases. Thus, having initially focused on the characteristics of the collective effort, including the stakeholders involved and the activities carried out, we began paying attention to the role of the orchestrators as well. Finally, through this iterative analytical process, we identified a third category of findings relevant to the research question, namely, how the combination of the characteristics of the collective effort and the orchestrators' role affected the outcome of the collective work. In short, the process of analysis proceeded from initial codes such as 'aim of initiative', 'stakeholders', 'activities', 'topics discussed', and 'outcome' to a broader categorisation of the findings into three themes: (1) collective activities and knowledge production, (2) the role of the initiators, and (3) the process' outcomes.

## Findings

In the following section, we present our findings, structured according to the three themes that emerged during the analysis: collective work for hybrid knowledge production, the role of the orchestrators, and the outcome of collective efforts.

### *Collective work for hybrid knowledge production*

The national inquiry process and the nationwide network brought together various stakeholders interested in promoting and facilitating the introduction of AI in Norwegian healthcare services. Both initiatives can be viewed as collective efforts, but with different modes of communication and interaction among participants.

In the case of the national inquiry process (case 1), the collective work, including knowledge sharing and production, among the involved stakeholders started with a series of six insight meetings, one for each group of stakeholders: 'Managers/project managers', 'Clinical staff', 'Procurers', 'User representatives and patient organisations', and 'Industry and research' (the latter group met twice). These meetings began with a presentation by the project team members from the Directorate of Health. They presented the rationale behind the process in relation to the policy plan (the NHSP), the coordination project, and finally, the course and aim of the inquiry process, highlighting the need for feedback from relevant experts (the participating stakeholders) on their needs and perceived

challenges regarding the introduction of AI in healthcare. After the presentations, the participants would then voice their concerns as the meetings proceeded into phases of discussion. For instance, in the meeting with the ‘industry and research’ stakeholders, when asked to point out what the premises to start using AI were, many concerns were listed. For instance, the following were mentioned:

*“Knowledge and trust must be built among healthcare professionals – [about what AI can do and cannot do] knowledge will create trust. We can’t just introduce a product without proving an effect. [...] we must have clear evidence that the solutions work – this is partly lacking in the products we see today. Relevance for our patients: Evidence is not necessarily the same as the technology being relevant because the solution can be based on other data than our own”* (observation notes, industry and research meeting).

This excerpt illustrates the openness of the initial round of inputs, spanning themes such as trust, AI effect, and clinical relevance.

Similarly, during the meetings with the procurement group, the participants’ input was open and varied, and they voiced concerns about various aspects of procuring AI solutions. For instance, they discussed what is particularly important to ensure good clinical practice when using AI, the perceived challenges for procurers, and how actors at the national level can contribute to addressing these issues. They also discussed how healthcare services could facilitate the use of AI, and whether the procurers needed guidance on procuring AI. As such, the participants shared their concerns, experiences, and understandings broadly. For instance, one of the participants expressed that:

*“If you ask what you can help me with: We use little AI today [in Norway], but we have started a thinking process: How should we take it into use? [...] What can an AI algorithm do for us? Should we put together such an algorithm that satisfies needs – the AI lab [developers] know what they can offer, the users do not necessarily know. Using helicopters as an example: We cannot fly helicopters from day one. We need to know something about turbulence, mechanics, etc.”* (Observation notes, procurement meeting).

Overall, these meetings were guided by a specific set of questions prepared by the project coordinators and opened up a range of themes in the form of concerns, hopes, and expectations. The insight meetings produced written summaries that included the stakeholders’ needs and challenges related to introducing AI in healthcare. During the meetings, participants shared knowledge not only with the coordinators but also among themselves, contributing collectively to the production of hybrid knowledge.

The insight meetings were followed by three workshops involving mixed stakeholder groups. The workshops aimed to discuss the findings from the insight meetings and generate ideas for further initiatives and possible solutions to enable AI in clinical practices. The topics discussed included access to sufficiently large high-quality data sets for research and product development, enhanced AI competencies among healthcare personnel, and guidance within current regulatory frameworks. They also discussed how to enable more collaboration between researchers, product developers, and clinicians, as well as ethical issues and healthcare personnel's information needs. Again, the outcome of these workshops was written summaries of ideas for solutions to the issues discussed.

Finally, various stakeholders were invited to two feedback meetings: one open to the public and announced on the project's website, and the other for invited people from radiology environments. The topics discussed aligned with those from the workshops. However, the aim was now to gather feedback on the opinions and suggested solutions raised during the workshops. As the process ended, findings from all meetings were processed into a report for the Ministry of Health and Care Services. The report included a list of recommended activities to be carried out in the following year(s), potentially by the public agencies involved in the project team.

In case 2, the nationwide network enabled collective knowledge production during its meetings and seminars. These events included presentations and/or panel discussions with time set aside for audience questions at the end. The presenters were typically invited experts from the research or hospital sector, the industry, the Ministry of Health and Care Services, various health authorities, and other relevant directorates or authorities (e.g., the data protection authority). Thus, unlike the inquiry process, the network meetings consisted mainly of presentations and discussions among specially invited participants rather than activities facilitating interaction and discussion among its members. It was through these meetings and presentations that the network attempted to achieve the main goal of enabling knowledge sharing between otherwise separate stakeholders, the importance of which was emphasized by one of the informants as follows:

*"It would be very wrong if people like me, pure technologists or statisticians or something like that, were to sit around working with artificial intelligence in the health service. It's no use. Then we'd get nowhere. So, having that dialogue across... and sharing expertise... is in a way a very basic prerequisite for us to achieve anything"*  
(Secretariat member, 4).

While the meeting agendas for the inquiry process were set by the organisers, the network's agendas were informed by a list of topics ranked by the members

through digital polls made available at the end of each meeting. The topics to vote on represented pressing issues that both the secretariat and the members perceived as important to address to enable AI in healthcare. This included topics such as 'Validation and adjustment of AI solutions to local conditions', 'Ethics and legislation regarding AI', 'Make or buy', 'Competence development within the sector', 'Policy and financial incentives', and 'Infrastructure and cyber security' (see Kannelønning, 2024a).

The fourth meeting in 2021 is an example of how the network's meetings were conducted and how various stakeholders were involved. This meeting, which lasted for four hours, focused on a topic voted on at the previous meeting: 'Introduction and testing of AI procurements – Practical experiences and the role of the authorities' (our translation). The meeting consisted of five presentations and a final open discussion. The presentations concerned 1) AI procurement based on the experiences from the first larger AI procurement process in Norwegian hospitals, 2) the final report of the national inquiry process (case 1) presented by the project manager, 3) how to proceed from AI development to AI products based on experiences from an AI development project, 4) the Norwegian Medical Products Agency (previous Norwegian Medicines Agency) on their role and tasks related to clinical testing of AI, and 5) the Norwegian Institute of Public Health's role and product portfolio in a new national system for introducing new methods. Most of these presentations were made available on the network's website.

### *The role of the orchestrators*

The national inquiry process and the nationwide network were collective efforts orchestrated in different ways.

In the national inquiry process (case 1), the collective effort was orchestrated by the initiating project team led by the Directorate of Health. They structured the process into insight meetings, workshops, and feedback meetings, and set the agenda for each meeting, including the modality of interaction (e.g., within or across stakeholder groups). In their communication to the invited participants, the team emphasised their openness to feedback and commentaries as this excerpt from the observation notes describes from the observer's perspective:

*"They [the organisers] keep emphasising that these [findings from the meetings regarding perceived needs/suggestions for further activities, etc.] are just suggestions. [They are saying:] you can give feedback if you think otherwise. We are very, very open to input. If you think we are doing something too badly in some way. Or in the wrong order. We want your voice to be heard"* (Observation notes).

By fostering an open feedback culture in meetings, they aimed to build consensus throughout the process. However, when the time came to produce the report, the organisers – representing specific authorities and public agencies – prioritised certain findings over others based on what they deemed most important.

The network (case 2), on the other hand, was only partly orchestrated by its secretariat. The secretariat members took the role of orchestrators by initiating the network and inviting colleagues and other stakeholders to participate. However, after this initial “recruitment”, people interested in joining could sign up and become members through the network’s website. Furthermore, the theme of the meetings was determined based on the members’ opinions, after which the secretariat prepared detailed agendas and invited presenters. In this way, the secretariat took on a stable role as organisers on behalf of the members rather than a dominant one steering the agenda and curating a specific outcome.

### *Outcome of collective efforts*

The activities of the inquiry process (case 1) were closely connected to the production and use of various documents. Apart from being initiated by a document (the NHSP), documents were produced before the stakeholder involvement began to inform stakeholders about the upcoming process, its background and aim, the organisation and course of activities. Documents were also used to communicate insights and feedback from the meetings back to the participants and other relevant stakeholders throughout the process. These documents, or meeting summaries, were typically shared via email and on the project’s website to give stakeholders the opportunity to provide feedback or comments if desired. The main findings were also used as points of departure for discussions in subsequent meetings. Thus, these documents served to inform, engage, and mobilise people. For instance, the feedback meetings were organised to get feedback on the summaries of findings from both the insight meetings and the workshops. An early version of the final document was also published on the website, allowing interested stakeholders to read and comment on it before finalisation.

The recommendations and suggestions for further activities prioritised in the final report concerned areas such as developing regulatory guidelines for AI procurements, offering courses on digital competence and basic AI understanding, stimulating to more interdisciplinary collaboration across the sector, improving the technical infrastructure (registries and databases) for easier data access/sharing, and establishing activities addressing ethical issues (Helsedirektoratet et al. 2021, 48–57). After its publication, the final report influenced further which

activities the agencies represented in the project team prioritised the following year(s). For instance, the report was sent to the Ministry of Health and Care Services, and through their yearly allocation letter to the Directorate of Health, they emphasised that the directorate was to continue focusing on adapting relevant national frameworks and prioritise supporting the healthcare services with guidance on legal issues related to AI:

*“The Directorate of Health are to continue the work of adapting national framework conditions so that the health and care services can safely use Artificial Intelligence for patient treatment. In 2022, especially interdisciplinary guidance related to legal issues is to be prioritised”* (Helse- og omsorgsdepartementet, 2022, p. 36, our translation).

However, as the final report addressed only a subset of the topics discussed by the stakeholders involved in the inquiry process, the prioritisation in both the report and the allocation letter similarly favoured a specific trajectory for further work, leaving unexplored other relevant issues raised by the collective. Consequently, crucial issues concerning and even contesting the future of AI in healthcare, as depicted in the NHSP’s vision, were not put on the agenda. This included worries about whether specific AI technologies are accurate enough to produce valuable outcomes for healthcare personnel, whether AI use in healthcare will cause a deskilling of physicians, whether the healthcare personnel have sufficient competencies to properly use AI, whether the personnel can come to rely too much on the AI result (automation bias), or whether (commercial) AI technologies are applicable for local Norwegian healthcare settings, including their specific patient populations.

The outcome of the network activities (case 2) was less evident or concrete, taking the form of knowledge sharing over time. However, this outcome was also affected by certain restraints. Despite gaining approximately 200 members, the network’s heterogeneity seemed somewhat limited. The list of members showed that most were from the hard sciences or the bureaucracy, reflecting a particular (technological) view of AI, while members representing AI users or AI technologies in real-world settings were hardly represented. As the topics perceived as most important to discuss by the participants concerned AI implementation and use in clinical practice, rather than, for instance, AI development or technical infrastructure, the overrepresentation of members with technical expertise, situated far from clinical practices, made it challenging to share experiences and produce the relevant knowledge. This included knowledge on topics such as how healthcare providers should validate AI technologies before

implementation, whether the outcome of AI use is fair and safe, independent of the patient group, whether the healthcare professionals as users can trust the accuracy of the AI results, or whether the hospitals are allowed (legally) to adjust commercial medical AI technologies to local conditions. Thus, these dimensions, similar to the unaddressed issues in case 1, remained unexplored uncertainties, though for slightly different reasons: Actors with the necessary experience and knowledge to enable the discussion were not fully represented in the network or in the sector more broadly. However, it was not only this skewed composition of members that inhibited inclusive knowledge-making. The focus and structure of the meetings and their digital form had an effect too. As the meetings mainly consisted of presentations by invited experts and panel discussions, the members became passive spectators except for the occasional opportunities to ask questions. At the same time, the digital meeting format created additional distance among the members and between them and the presenters. Digital meetings offer fewer opportunities for informal or spontaneous interactions, as, for instance, in-person coffee breaks might do. They also make it harder to interrupt or claim space to speak during the meetings (Lampa, Sonnentheil, Tökés, & Warner, 2021, p. 4). Other constraining factors for interaction in digital meetings include the ability to turn the camera on and off at will, or the ease with which participants can distract themselves by working on other tasks on their computers while seemingly listening in.

Although the network's knowledge production was less tangible (no summaries or reports were produced), a secondary outcome was achieved through the secretariat's involvement in various activities outside the network, acting on its behalf. This included participating in workshops organised by the coordination project in case 1 on AI adoption in healthcare, writing a letter for the Ministry of Health and Care Services emphasising the need to pay more attention to the process of enabling AI in healthcare, and hosting an event at Arendalsuka in 2022, Norway's largest annual political gathering. On this occasion, they highlighted the need for a national "roadmap" to address the challenges of introducing AI in healthcare services.

## Discussion

Nordic countries are currently introducing AI into public healthcare services, and these initiatives reflect their tradition of policy-making and societal development characterised by consensus-building, dialogue, inclusion, and collective action

(Kuhnle & Alestalo, 2017; Lægneid, 2017). While research indicates that a collective dimension is crucial for ensuring sustainable and ethical outcomes of AI implementation in society (Floridi et al., 2018), there is still limited understanding of how collective efforts contribute to the introduction of AI. In this paper, we have analysed two experiences of collective efforts in the Norwegian healthcare sector while addressing the following research question: *How can collective efforts be organised and carried out to contribute to a consensus-based and inclusive introduction of AI in healthcare?*

In line with literature on collective efforts (e.g., Olson, 1965; Hardin, 1968), our findings suggest that processes of mobilising heterogeneous actors for collective agency are complex and can take different forms with different implications. The collective efforts described in the two cases of the present study both involve a large number of stakeholders with diverse expertise and practices. Their main activity is meetings, where knowledge is negotiated, shared, and generated as various topics are presented and discussed. However, the collective efforts of the two cases differed in types of initiators, orchestration modes, activities, engagement of the participants, duration, goals, and outcomes. In the following subsections, we discuss the implications of such differences in collective efforts in relation to hybrid forums and their orchestration, hybrid knowledge, and AI enactment.

### *Collective approaches and orchestration mode*

Our findings suggest that the national inquiry process (case 1) is a collective effort formally established as a “top-down” initiative with an orchestration mode predominantly resembling the dominating mode (Reypens et al., 2021). The project was initiated by a policy plan, and the organisers set the agenda and selected which parts of the produced knowledge was included in the final report. Still, during the meetings with the stakeholders, we also observed a consensus-based mode at work. The meetings were forums for open discussion, with the main aim of facilitating the sharing of knowledge, expertise, and opinions among 100 participants representing a wide variety of practices and realities. However, the organisers’ dominant role made visible the limits of the consensus-based mode, which was confined to meetings and had limited effect beyond these. In a way, the politics of the collective efforts make the consensus-based mode function as a legitimisation for the higher-level political agendas of the organisers, thereby limiting the collective character of the effort.

In the second case, variations in orchestration modes were also observed, but these were less pronounced. Even though the orchestration of the network had characteristics like the dominating mode – the secretariat initiated the network, formulated the aim, and organised the activities – many of its characteristics also aligned with the consensus-oriented orchestration mode. The network was an informally organised “bottom-up” initiative open to anyone interested, with an agenda set by the participants through voting. Our findings suggest that the members’ interests were considered and, to a certain extent, met, at least with respect to setting up the meeting agendas. The orchestrators also aimed to develop a common understanding of the voted topics by inviting presenters with relevant expertise. Thus, in this case, the dominating role mainly worked by facilitating activity coordination, while the consensus-based mode at work in the meetings had effects beyond the meetings, shaping the agendas for what would be further discussed. The two orchestrators’ modes were more aligned in this case than in the first.

These findings are in line with Reypens et al. (2021), showing that collective efforts can be orchestrated according to various modalities. Our findings show that it is not only the co-existence of different modalities that matter, but also their relationship. This relationship shapes how the collective effort develops over time, the knowledge it produces, and what outcome it achieves.

### *Hybrid knowledge*

Overall, our findings suggest that the collective approach in both cases contributed to the introduction of AI in healthcare by creating a meeting place and a knowledge-sharing opportunity across expertise and domains of knowledge. In both initiatives, hybrid knowledge was produced to varying degrees through guided discussions and presentations. This knowledge was further concretised in official and internal reports, presentation slides, meeting minutes, meeting invitations, and letters. However, the production of hybrid knowledge faced limitations in both cases.

In the first case, the inquiry process, the production of hybrid knowledge was constrained by participation and the role of the orchestrator. The number of stakeholders involved at each meeting was never especially high, and the expertise, at least at the initial group-based insight meetings, was not particularly diverse or hybrid. This aspect affected the outcome of the process as reflected in the final report. In addition, as the report was carefully curated, and crucial issues raised by the participants remained unaddressed and further unexplored, hybrid

knowledge was only partially produced. As the report was mainly based on enactments of AI introduction as technically and legally challenging, the follow-up work prioritised by the agencies of the project team was subsequently steered in those two directions – technical and legal.

In the second case, our findings suggest that despite the network's open, flexible, and inclusive approach, its ability to produce hybrid knowledge was challenged by the very topics members perceived as most crucial for enabling AI implementations. These topics were the most voted-for discussion points, remaining at the top of the list over the two years of observation. Again, this limit was due to an overrepresentation of technology-focused interests, while (potential) AI users, representing the practices AI was meant to become a part of, were scarce in the member list. Consequently, the knowledge produced was not particularly hybrid, which in turn influenced the type of knowledge shared as the outcome of the network's collective work. Other constraining factors for the network's hybrid knowledge production were the limited opportunities for actual interaction and knowledge sharing among the participants during the meetings, only partly occasioned by the meeting structure (mainly presentations) and their digital format.

Based on Olson's (1965) argument that the larger and more heterogeneous a group is, the harder it becomes for participants to see their interests served, such passive attendance in a large network might lead to a loss of interest and a further decline in contributions from actors who become merely spectators. Therefore, despite Ostrom's (2005) claim that larger self-organised collectives could overcome this "tragedy of the commons", the lack of interaction constitutes a limiting factor for hybrid knowledge production.

### *Hybrid forums and the politics of AI enactment*

The discussions unfolding in the two cases and their outputs contributed to moving the introduction of AI in healthcare one step forward in certain directions, even if not directly materialising in actual implementations. In line with the purpose of hybrid forums as outlined by Gibbons et al. (1994) and Callon et al. (2011), the study's findings further suggest that the collective approaches contributed to a secondary outcome in terms of a maturing process for the participants and the orchestrators in their understanding of the complexities, opportunities, and challenges of introducing AI in healthcare and to the production of shared awareness within the collectives.

However, the findings also show that, in both cases, hybrid knowledge is not just about combining perspectives and visions into a shared understanding through negotiation, but also about enacting conflicting visions of AI in healthcare, where some are more dominant and thus more visible than others. In both cases of the present study, certain technology-focused futures with AI in healthcare were enacted, which could potentially clash with the realities of less-represented practices, particularly as AI is being implemented (Mol, 2002; Suchman, 2007).

Thus, overall, our findings show that in hybrid forums, where actors such as technologists, clinicians, hospital managers, and health authorities are gathered, different “AI technologies in healthcare” are enacted depending on who the actors are, what they voice, and what they consider as important priorities in relation to their knowledge. As such, the stakeholders within hybrid forums represent various, partially connected sets of realities, and as their enactments interfere with one another, the different practices develop. Based on this, decisions concerning what is crucial for introducing AI in healthcare, which kind of AI should be implemented, how, and for whom, are also political (Mol, 2002). As our findings show, the practices of the most represented and powerful actors within the hybrid forums will strongly influence how the future of AI in healthcare unfolds.

### *Implications for designing future collective initiatives*

Building on our findings, we draw three implications for future initiatives that aim to bring together heterogeneous stakeholders to gain knowledge and a deeper understanding of the challenges associated with the introduction of complex emerging technologies, such as AI, into society.

First, we suggest that the work of including stakeholders as participants plays a critical role. The participants bring their expertise and knowledge to the collective discussions. These discussions and the knowledge produced as they unfold are, in turn, shaped by both the present and the absent expertise and knowledge. As in the case of the network of professionals, the discussion did not move forward on topics such as validation processes, ethical concerns, and the (mis)use of AI, in part due to limited hybridity in the expertise represented among the participants and other contributors. Based on this, we argue that striving to include actors representing multiple practices and to facilitate ways to explore issues related to missing knowledge should be prioritised when aiming for hybrid knowledge production. In this context, paying attention to further potentially constraining factors, such as less inclusive or less engaging meeting formats (digital or not), could also be crucial.

Second, our findings suggest that agenda-setting and the curation of the topics discussed and reported through collective efforts also play a critical role. Here, we recognise and confirm the critical role of orchestrators. For instance, in the case of the national inquiry process (case 1), the project team selected and guided both the discussion and the final reporting. This work was shaped by the national political agenda and a strong vision of the role of AI technologies in transforming healthcare services. Exploring, rather than overlooking, emerging uncertainties that challenge the dominant technological vision could further increase the likelihood of a successful implementation process, as knowledge of potential future challenges can already be addressed or mitigated.

Finally, this paper contributes to research interested in understanding how advanced and complex technologies, such as AI, are introduced in healthcare services through collective efforts. Recent research has shown that heterogeneous constellations of actors are mobilised in different ways (Kannelønning, 2024b) and how controversies emerge through such collectivity, revealing crucial challenges and uncertainties in the introduction of AI (Kannelønning, 2023). Research has also recognised that many aspects and applications of AI will require some form of collective action, meaning that collectives should be involved in improving the outcomes of AI development to ensure that AI contributes to good outcomes for society (de Neufville & Baum, 2021; Floridi et al., 2018). We find this research to be especially relevant in the context of the Nordic countries and their experiences with collective participation and consensus- and dialogue-oriented policy traditions.

## Final remarks

The results of the longitudinal case study presented in this paper suggest that in consensus-oriented countries like the Nordics, establishing collective efforts can be a way to address and improve understanding of pressing challenges associated with the introduction of emerging advanced technologies. Initiatives involving a broad set of stakeholders can also use such collective efforts to legitimise the need for certain solutions or future activities to be carried out by, for instance, the organisers themselves, regardless of the actual knowledge produced by the collective. However, when the technology in focus moves towards implementation, the overlooked or not yet fully explored challenges and uncertainties brought to the surface within the hybrid forums might strike back and require (costly) adjustments or a resetting of the course of the technology introduction. Thus, aiming for multiplicity both among the actors involved in the process and in the

accumulated outcome could be key to preventing collective efforts from becoming purely cosmetic.

Thus, we suggest that in future initiatives aimed at fostering collective efforts focused on consensus-building to better understand and cope with challenges in technology introductions, a hybrid orchestration mode can be strategically applied across various phases of the project. However, these initiatives should pay attention to how and when different orchestration modes are enabled to ensure that they do not undermine the aim and strength of collective approaches to sustainable futures with the emerging technology in focus. If collective efforts are not achieved or if their mission to produce hybrid knowledge somehow fails, the risk might be, as de Neufville and Baum argue, that “society may stumble blindly into outcomes that are collectively worse” (de Neufville & Baum, 2021, p. 9).

## References

- Alami, H., Lehoux, P., Auclair, Y., Guise, M. d., Gagnon, M.-P., Shaw, J., . . . Ag Ahmed, M. A. (2020). Artificial Intelligence and Health Technology Assessment: Anticipating a New Level of Complexity. *Journal of Medical Internet Research*, 22(7), 1-12. doi:10.2196/17707
- Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A., Almohareb, S. N., . . . Albekairy, A. (2023). Revolutionizing healthcare: the role of artificial intelligence in clinical practice. *BMC Medical Education*, 23(1). doi:10.1186/s12909-023-04698-z
- Callon, M., Lascoumes, P., & Barthe, Y. (2011). *Acting in an uncertain world: An essay on technical democracy*. MIT press.
- Danish Ministry of Finance, & Danish Ministry of Industry, B. a. F. A. (2019). *National Strategy for Artificial Intelligence*. Retrieved from [https://en.digst.dk/media/19337/305755\\_gb\\_version\\_final-a.pdf](https://en.digst.dk/media/19337/305755_gb_version_final-a.pdf)
- de Laet, M., & Mol, A. (2000). The Zimbabwe Bush Pump: Mechanics of a Fluid Technology. *Social Studies of Science*, 30(2), 225-263. doi:10.1177/030631200030002002
- de Neufville, R., & Baum, S. D. (2021). Collective action on artificial intelligence: A primer and review. *Technology in Society*, 66, 101649. doi:10.1016/j.techsoc.2021.101649
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., . . . Rossi, F. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707. doi:10.1007/s11023-018-9482-5
- Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P., & Trow, M.

(1994). *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. London: Sage Publications, Limited.

Hardin, G. (1968). The Tragedy of the Commons. *Science*, 162(3859), 1243-1248. doi:10.1126/science.162.3859.1243

Helse- og omsorgsdepartementet. (2019). *Nasjonal helse- og sykehusplan 2020-2023*. (Meld. St. 7 (2019–2020)). Retrieved from <https://www.regjeringen.no/no/dokumenter/meld.-st.-7-20192020/id2678667/>

Meld. St. 7 (2019–2020). *National Health and Hospital Plan 2020-2023. Summary*. Retrieved from <https://www.regjeringen.no/contentassets/95eec808f0434acf942fca449ca35386/en-gb/pdfs/stm201920200007000engpdfs.pdf>

Helse- og omsorgsdepartementet. (2022). *Tildelingsbrev til Helsedirektoratet for 2022*. Retrieved from <https://www.regjeringen.no/globalassets/departementene/hod/tildeling-oppdrag-og-arsrapporter/2022/tildelingsbrev-til-helsedirektoratet-for-2022.pdf>

Helse- og omsorgsdepartementet. (2024). Meld. St. 9 (2023–2024). *Nasjonal helse- og samhandlingsplan 2024–2027*. Retrieved from <https://www.regjeringen.no/contentassets/4e5d9e6c63d24cd7bdab5d8c58d8adc4/no/pdfs/stm202320240009000dddpdfs.pdf>

Helsedirektoratet, Direktoratet for e-helse, Helsetilsynet, & legemiddelverk, S. (2021). *Tilrettelegging for bruk av kunstig intelligens i helsetjenesten. Med utgangspunkt i det radiologiske fagområdet*. Retrieved from Oslo: <https://www.helsedirektoratet.no/rapporter/tilrettelegging-for-bruk-av-kunstig-intelligens-i-helsesektoren-ny-01.10.2021/>

Kannelønning, M. S. (2023). Contesting futures of Artificial Intelligence (AI) in healthcare: formal expectations meet informal anticipations. *Technology Analysis & Strategic Management*, 1-12. doi:10.1080/09537325.2023.2226243

Kannelønning, M. S. (2024a). Navigating uncertainties of introducing artificial intelligence (AI) in healthcare: The role of a Norwegian network of professionals. *Technology in society*, 76, 102432. doi:10.1016/j.techsoc.2023.102432

Kannelønning, M. S. (2024b). 'How will it hit us?': introducing artificial intelligence (AI) in the Norwegian healthcare services : three modes of actor mobilisation OsloMet - Oslo Metropolitan University. Oslo. <https://hdl.handle.net/11250/3112563>

<https://hdl.handle.net/11250/3112563>

Kuhnle, S., & Alestalo, M. (2017). The modern Scandinavian welfare state. In P. Nedergaard & A. Wivel (Eds.), *The Routledge Handbook of Scandinavian Politics*. London: Routledge.

Læg Reid, P. (2017). Nordic administrative traditions. In P. Nedergaard & A. Wivel (Eds.), *The Routledge Handbook of Scandinavian Politics*. London: Routledge.

Lampa, E., Sonnentheil, B., Tökés, A., & Warner, G. (2021). What has the COVID-19 pandemic taught us about conducting patient and public involvement remotely? Insights from a series of digital meeting observations. *Research Involvement and Engagement*, 7(1), 1-8. doi:10.1186/s40900-021-00315-9

Mol, A. (2002). *The Body Multiple: Ontology in Medical Practice*. Durham: Duke University Press.

Noorbakhsh-Sabet, N., Zand, R., Zhang, Y., & Abedi, V. (2019). Artificial Intelligence Transforms the Future of Health Care. *The American Journal of Medicine*, 132(7), 795-801. doi:10.1016/j.amjmed.2019.01.017

Olson, M. (1965). *The logic of collective action: public goods and the theory of groups* (Vol. 124). Cambridge, Mass: Harvard University Press.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.

Ostrom, E. (2005). *Understanding Institutional Diversity*. New Jersey: Princeton University Press.

Periáñez, Á., Fernández Del Río, A., Nazarov, I., Jané, E., Hassan, M., Rastogi, A., & Tang, D. (2024). The Digital Transformation in Health: How AI Can Improve the Performance of Health Systems. *Health Systems & Reform*, 10(2), 2387138. doi:10.1080/23288604.2024.2387138

Reypens, C., Lievens, A., & Blazevic, V. (2021). Hybrid Orchestration in Multi-stakeholder Innovation Networks: Practices of mobilizing multiple, diverse stakeholders across organizational boundaries. *Organization Studies*, 42(1), 61-83. doi:10.1177/0170840619868268

Suchman, L. (2007). *Human-Machine Reconfigurations: Plans and Situated Actions*. Cambridge: Cambridge University Press.

Timmermans, S., & Tavori, I. (2012). Theory Construction in Qualitative Research: From Grounded Theory to Abductive Analysis. *Sociological Theory*, 30(3), 167-186. doi:10.1177/0735275112457914

# Between Theoretical and Empirical Ethics of Healthcare Artificial Intelligence

## The Case of Autonomy in Breast Cancer Screening

Victor Vadmand Jensen<sup>1,2</sup>

Jens Christian Bjerring<sup>3</sup>

---

<sup>1</sup>Institut for Klinisk Medicin, Aarhus Universitet

<sup>2</sup>Universitetsklinik for Interdisciplinære Ortopædkirurgiske Forløb, Regionshospitalet Silkeborg

<sup>3</sup>Institut for Kultur og Samfund, Aarhus Universitet

*vvoj@clin.au.dk*

*filjcb@cas.au.dk*

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

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

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

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

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

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

### Introduction

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

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

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

*"Although participants were open to the use of novel AI solutions in their care, most of them, especially those with diabetes, would prefer for them to be implemented alongside clinicians to keep human control and doctor-patient relationship in medicine" (Schaarup et al., 2023, p. 6).*

*"These results highlight a potential concern that AI will impact the therapeutic alliance and potentially reduce the positive relational elements (i.e. empathy, understanding, etc.) present within the doctor-as-a-person elements of patient-centered care" (Witkowski et al., 2024, p. 8).*

*"Many respondents at some point entertained the thought of the AI, at some point, being able to completely automate the tasks of the physicians, but always immediately observed that this would be a very bad evolution" (Van Cauwenberge et al., 2022, p. 8).*

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

We can think of such ethical stances as establishing broad, overarching ideals that are intended to apply universally across clinical practices while deliberately avoiding a case-by-case focus on individual circumstances (Beauchamp & Childress, 2001). This approach typically understands ethical principles *theoretically*: as principles that can be reasoned about abstractly rather than derived from empirical realities. As Bærøe et al. (2017, p. 7) argue, the aim of theoretical ethics is to develop "general ethical principles for medicine and health

promotion” without direct reference to specific human circumstances. Principles of theoretical ethics impose non-negotiable requirements, as seen, for example, in AI governance frameworks (AI HLEG, 2019), where conditions for AI use are set independently of potential benefits that might arise in particular cases (Haward & Janvier, 2019).

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

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

First, we survey research and policy concerns about autonomous AI systems in healthcare and relate them to familiar autonomy ideals in biomedical ethics, distilling these into the Non-Autonomous Principle (Section 2). Second, drawing on policy and press documents, we examine Danish breast cancer screening, where AI is increasingly deployed as first reader of mammograms (Section 3). This setting is relevant for testing the Non-Autonomous AI Principle because, in practice, the

first reader AI has an authority comparable to a radiologist: its assessments carry equal weight and cannot be unilaterally overruled by the human second reader. Moreover, since this use of healthcare AI is by now common in Denmark, it hence also carries a seeming ethical legitimacy from its routine use in clinical practice and institutional promotion.

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

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

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

Second, we do not assume that “empirical ethics” must oppose strict rules or principles. Empirically informed approaches can yield relatively rigid norms. The difference is that these norms are reconstructed from within practices rather than

both stipulated prior to and defended largely independently of implementation and negotiation in concrete settings. Likewise, theoretical approaches can, and often do, take empirical realities seriously. But their starting point typically lies elsewhere: They begin with principles, rules, or ideals which are then applied to practice. In what follows, we use “theoretical” and “empirical” as ideal-typical starting points rather than as mutually exclusive camps. The former is principle-first, abstracting from practice; the latter is practice-first, grounded in intra-normativity. We treat the Non-Autonomous AI Principle as a prominent example of the former and our core claim is that an intra-normative, practice-first analysis of breast cancer screening explains why its assumed universality should be questioned.

## The Non-Autonomous AI Principle as Foundational for Healthcare AI

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

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

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

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

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

### *Respect for Autonomy*

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

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

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

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

### *Human Oversight*

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

The first reason relates to errors and bias. Because AI systems learn statistical patterns from training data, they will inevitably make errors in their classifications – even in settings that look routine. This risk is amplified when no physician reviews the system's outputs (Filippi et al., 2023; Najjar, 2023). Unlike physicians, AI does not reflect on and correct its own mistakes during its decision-making which can cause frustration among healthcare personnel (Beede et al., 2020). Evidence also indicates that physicians are more forgiving of human errors than AI errors (Lenskjold et al., 2023). At the same time, reliance on AI may introduce

automation bias, nudging clinicians to over-trust AI-generated outputs (Bernstein et al., 2023; Fang et al., 2024). In AI-assisted breast cancer screening, this may for instance lead to unnecessary and costly patient recalls (Hernström et al., 2025).

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

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

### *No Final Authority*

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

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

### *The Non-Autonomous AI Principle*

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

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

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

collectively, presenting them as if they formed a single, unified constraint on acceptable AI use.

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

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

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

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

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

## *The Case – AI as First Reader in Breast Cancer Screening*

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

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

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

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

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

### *An AI First Reader and the Non-Autonomous AI Principle*

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

This practice, we suggest, challenges each element of the Non-Autonomous AI Principle. For the first condition concerning respect for autonomy, note that the radiologists’ professional judgment at the first reading stage is structurally constrained: they are required to act on disagreement with the AI first reader. That is, when it comes to radiologist autonomy, the second reader is structurally required to treat the AI’s assessment as an independent judgment containing recognized institutional consequences, i.e. as triggering a consensus conference upon disagreement. This workflow thus prevents the radiologist from simply treating the AI system as an imperfect tool that can be ignored in a specific case. Professional discretion gets limited by procedure rather than by medical judgment

alone. Thus, the physicians' ability to act on their values and professional judgment is partly displaced by a non-answerable, non-responsible system and the rules built around it. Likewise, when it comes to patient autonomy, patients encounter screening as a highly standardized procedure with little room for expressing individual values. Since the introduction of the relevant AI system is explicitly justified in terms of efficiency and speed, its inclusion does not speak to any richer forms of patient self-determination.

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

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

Accordingly, given the use of AI in breast cancer screening is regarded as ethically acceptable – which current practice clearly supports – these observations cast doubt on the Non-Autonomous AI Principle's universal applicability. Concretely, autonomy is reshaped by binding both patients and physicians into highly standardized AI-driven routines; meaningful human oversight is minimized at the level of individual cases; and AI's application in the workflow functions as a decision-maker and gatekeeper. Of course, not every technical substitution for

a human function diminishes autonomy or responsibility. In many cases, like the use of routine measuring instruments, functional substitution leaves the underlying normative structure largely intact. What the breast cancer screening case shows, however, is that *this particular implementation* of AI places significant pressure on all three elements of the Non-Autonomous AI Principle – above all, the prohibition on AI occupying roles with decisional power comparable to that of human.

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

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

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

# The Intra-normativity of Autonomy

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

## *Structural Issues*

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

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

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

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

### *Existential Concerns*

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

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

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

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

These studies suggest that moving through cancer care testing, diagnosis, treatment, and follow-up triggers emotional turmoil. Frumer et al. (2021) describe

how patients may experience this process as a kind of suspended state, “being in the meantime” (Frumer et al., 2021, p. 19). This meantime is filled with ambiguities shaping how both patients and their relatives structure their lives. Uncertainty about a potential diagnosis can trigger both fear and apathy in patients and relatives (Frumer et al., 2021). Moreover, receiving a cancer diagnosis is an unsettling experience, with patients reporting how “[cancer, ed.] and the prospect of extensive treatment evoked a profound fear, leading to existential distress” (Wilhøft Kristensen et al., 2024, p. 3).

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

## Concluding Remarks

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

Considered through empirically oriented ethics, the implications become clear. Rather than relying solely on general, principle-first accounts, we also need a

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

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

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

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

By aligning AI systems with the concrete norms, care need, and workflow pressures of specific clinical contexts, we can thus design ethically acceptable AI

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

## References

- Adams, S. J., Tang, R., & Babyn, P. (2020). Patient Perspectives and Priorities Regarding Artificial Intelligence in Radiology: Opportunities for Patient-Centered Radiology. *Journal of the American College of Radiology*, 17(8), 1034–1036. <https://doi.org/10.1016/j.jacr.2020.01.007>
- Aglund, M. (2024, July 6). *Angst ved en kræftsygdom—Kræftens Bekæmpelse*. Kræftens Bekæmpelse. <https://www.cancer.dk/hverdag/dit-liv-med-kræft/reaktioner-og-tanker/angst/>
- AI HLEG. (2019). *ETHICS GUIDELINES FOR TRUSTWORTHY AI*. European Commission. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- Bærøe, K., & Bringedal, B. (2022). Social inequalities in health: The need for a new concept of health. *Tidsskrift for Forskning i Sygdom Og Samfund*, 19(36), 27–38. <https://doi.org/10.7146/tfss.v19i36.133048>
- Bærøe, K., Ives, J., De Vries, M., & Schildmann, J. (2017). On classifying the field of medical ethics. *BMC Medical Ethics*, 18(1), 30, s12910-017-0193–x. <https://doi.org/10.1186/s12910-017-0193-x>
- Beauchamp, T. L., & Childress, J. F. (2001). *Principles of Biomedical Ethics* (5th ed.). Oxford University Press. <https://ebookcentral-proquest-com.ez.statsbiblioteket.dk/lib/asb/detail.action?docID=5763592>
- Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., & Vardoulakis, L. M. (2020). A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12. <https://doi.org/10.1145/3313831.3376718>
- Bergquist, M., Rolandsson, B., Gryska, E., Laesser, M., Hoefling, N., Heckemann, R., Schneiderman, J. F., & Björkman-Burtscher, I. M. (2023). Trust and stakeholder perspectives on the implementation of AI tools in clinical radiology. *European Radiology*, 34(1), 338–347. <https://doi.org/10.1007/s00330-023-09967-5>
- Bernstein, M. H., Atalay, M. K., Dibble, E. H., Maxwell, A. W. P., Karam, A. R., Agarwal, S., Ward, R. C., Healey, T. T., & Baird, G. L. (2023). Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. *European Radiology*, 33(11), 8263–8269. <https://doi.org/10.1007/s00330-023-09747-1>
- Bjerring, J. C., Busch, J., & Aastrup Munch, L. (2025). A Counterfactual Account of Algorithmic Robustness. *Minds and Machines*, 35(3), 34. <https://doi.org/10.1007/s11023-025-09734-z>

- Blomberg, S. N., Christensen, H. C., Lippert, F., Ersbøll, A. K., Torp-Petersen, C., Sayre, M. R., Kudenchuk, P. J., & Folke, F. (2021). Effect of Machine Learning on Dispatcher Recognition of Out-of-Hospital Cardiac Arrest During Calls to Emergency Medical Services: A Randomized Clinical Trial. *JAMA Network Open*, 4(1), e2032320. <https://doi.org/10.1001/jamanetworkopen.2020.32320>
- Bugge, M. (2023, July 16). *Hovedstaden får kunstig intelligens til at opdage brystkræft—Resten af landet vil gøre den kunsten efter*. DR. <https://www.dr.dk/nyheder/viden/teknologi/hovedstaden-faar-kunstig-intelligens-til-opdage-brystkraeft-resten-af-landet>
- Danish Health Authority. (n.d.). *Information about the Danish breast cancer screening programme (Health for All)*. Danish Health Authority. Retrieved September 15, 2024, from <https://www.sst.dk/-/media/Udgivelser/2023/Kraeft/Screening/UK-Pjece-brystkraeft.ashx>
- Danish Health Authority. (2023, October 10). *Screening for brystkræft*. Sundhedsstyrelsen. <http://www.sst.dk/da/Borger/Sygdomme-og-lidelser/Kræftsygdom/Screening-for-kraeft/Brystkraeft>
- Danske Regioner. (2024). *Vi skal skabe et godt og effektivt sundhedsvæsen ved brug af kunstig intelligens*. Danske Regioner. <https://www.regioner.dk/media/0urf11dv/kunstig-intelligens-udspil-070624.pdf>
- Djursing, T. (2017, October 23). *Doktor Watson i modvind: Foreslog livsfarlig medicin til danske patienter | Ingeniøren*. Ingeniøren. <https://ing.dk/artikel/doktor-watson-i-modvind-foreslog-livsfarlig-medicin-til-danske-patienter>
- Elhakim, M. T., Graumann, O., Larsen, L. B., Nielsen, M., & Rasmussen, B. S. (2020). Kunstig intelligens til cancerdiagnostik i brystkræftscreening. *Ugeskr Læger*. <https://ugeskriftet.dk/videnskab/kunstig-intelligens-til-cancerdiagnostik-i-brystkraeftscreening>
- European Commission. (2021). *Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS*. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- Fajstrup, M. (2021, July 26). *AI og visuel feedback ruster lægerne til vanskelig hudkræftdiagnose | Version2*. <https://www.version2.dk/artikel/ai-og-visuel-feedback-ruster-laegerne-til-vanskelig-hudkraeftdiagnose>
- Fang, N., Pluyter, J., Bakker, S., Jacobs, I., Luyer, M., Nederend, J., Raijmakers, J., Chen, L.-L., & Funk, M. (2024). From Experience to Experience: Key Insights for Improved Interaction with AI in Radiology. *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 1–7. <https://doi.org/10.1145/3613905.3650779>
- Filippi, C. G., Stein, J. M., Wang, Z., Bakas, S., Liu, Y., Chang, P. D., Lui, Y., Hess, C., Barboriak, D. P., Flanders, A. E., Wintermark, M., Zaharchuk, G., & Wu, O. (2023). Ethical Considerations and Fairness in the Use of Artificial Intelligence for Neuroradiology. *American Journal of Neuroradiology*, 44(11), 1242–1248. <https://doi.org/10.3174/ajnr.A7963>
- Frumer, M., Andersen, R. S., Vedsted, P., & Offersen, S. M. H. (2021). 'In the Meantime': Ordinary Life in Continuous Medical Testing for Lung Cancer. *Medicine Anthropology Theory*, 8(2), 1–26. <https://doi.org/10.17157/mat.8.2.5085>
- Goh, E., Gallo, R., Hom, J., Strong, E., Weng, Y., Kerman, H., Cool, J. A., Kanjee, Z., Parsons, A. S., Ahuja, N., Horvitz, E., Yang, D., Milstein, A., Olson, A. P. J., Rodman, A., & Chen, J. H. (2024). Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial. *JAMA Network Open*, 7(10), e2440969. <https://doi.org/10.1001/jamanetworkopen.2024.40969>

- Haward, M. F., & Janvier, A. (2019). Empirical Over Theoretical Ethics: Choosing What Matters to Patients and Families. *The American Journal of Bioethics*, 19(3), 54–56. <https://doi.org/10.1080/15265161.2019.1572813>
- Hernström, V., Josefsson, V., Sartor, H., Schmidt, D., Larsson, A.-M., Hofvind, S., Andersson, I., Rosso, A., Hagberg, O., & Lång, K. (2025). Screening performance and characteristics of breast cancer detected in the Mammography Screening with Artificial Intelligence trial (MASAI): A randomised, controlled, parallel-group, non-inferiority, single-blinded, screening accuracy study. *The Lancet Digital Health*, S258975002400267X. [https://doi.org/10.1016/S2589-7500\(24\)00267-X](https://doi.org/10.1016/S2589-7500(24)00267-X)
- Hudnall, C. (2021, June 1). *Thinking About AI?* American College of Radiology. <https://www.acr.org/Practice-Management-Quality-Informatics/ACR-Bulletin/Articles/June-2021/Thinking-About-AI>
- Hyldgaard, S. (2024, April 2). *Kunstig intelligens opsporer brystkræft og afhjælper ressourcemangel i Region Nordjylland*. <https://rn.dk/Service/Nyhedsbase-RN/Nyhed?id={625E1BB4-3512-410C-8023-B7E5136518E3}>
- Jenvall, L. (2024, June 1). *Ny økonomaftale giver store smil, men regionerne er fanget i et krydspres*. DR. <https://www.dr.dk/nyheder/politik/ny-oekonomiaftale-giver-store-smil-men-regionerne-er-fanget-i-et-krydspres>
- Kambhamettu, H., Metaxa, D., Johnson, K., & Head, A. (2024). Explainable Notes: Examining How to Unlock Meaning in Medical Notes with Interactivity and Artificial Intelligence. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–19. <https://doi.org/10.1145/3613904.3642573>
- Kommission for robusthed i sundhedsvæsenet. (n.d.). *Robusthedskommissionens anbefalinger—Bilagsrapport*. Retrieved December 18, 2024, from <https://www.ism.dk/Media/638336462433905130/Robusthed-Bilagsrapport-TILG.pdf>
- Kühl, J., Elhakim, M. T., Stougaard, S. W., Rasmussen, B. S. B., Nielsen, M., Gerke, O., Larsen, L. B., & Graumann, O. (2023). Population-wide evaluation of artificial intelligence and radiologist assessment of screening mammograms. *European Radiology*, 34, 3935–3946. <https://doi.org/10.1007/s00330-023-10423-7>
- Lenskjold, A., Nybing, J. U., Trampedach, C., Galsgaard, A., Brejnebøl, M. W., Raaschou, H., Rose, M. H., & Boesen, M. (2023). Should artificial intelligence have lower acceptable error rates than humans? *BJR|Open*, 5(1), 20220053. <https://doi.org/10.1259/bjro.20220053>
- Luchini, C., Pea, A., & Scarpa, A. (2022). Artificial intelligence in oncology: Current applications and future perspectives. *British Journal of Cancer*, 126(1), 4–9. <https://doi.org/10.1038/s41416-021-01633-1>
- Mack, J. W., Block, S. D., Nilsson, M., Wright, A., Trice, E., Friedlander, R., Paulk, E., & Prigerson, H. G. (2009). Measuring therapeutic alliance between oncologists and patients with advanced cancer: The Human Connection Scale. *Cancer*, 115(14), 3302–3311. <https://doi.org/10.1002/cncr.24360>
- Mainz, J., Munch, L., Bjerring, J. C., & Godtfredsen, S. (2023). Why algorithmic speed can be more important than algorithmic accuracy. *Clinical Ethics*, 18(2), 161–164. <https://doi.org/10.1177/14777509221138750>
- Maris, M. T., Koçar, A., Willems, D. L., Pols, J., Tan, H. L., Lindinger, G. L., & Bak, M. A. R. (2024). Ethical use of artificial intelligence to prevent sudden cardiac death: An interview study of patient perspectives. *BMC Medical Ethics*, 25(1), 42. <https://doi.org/10.1186/s12910-024-01042-y>

- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>
- Mudgal, K. S., & Das, N. (2024). The ethical adoption of artificial intelligence in radiology. *BJR|Open*, 2(1), 20190020. <https://doi.org/10.1259/bjro.20190020>
- Najjar, R. (2023). Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging. *Diagnostics*, 13(17), 2760. <https://doi.org/10.3390/diagnostics13172760>
- Offersen, S. M. H., Risør, M. B., Vedsted, P., & Andersen, R. S. (2018). Cancer-before-cancer. *Medicine Anthropology Theory*, 5(5). <https://doi.org/10.17157/mat.5.5.540>
- Pellegrino, E. D. (1994). Patient and physician autonomy: Conflicting rights and obligations in the physician-patient relationship. *The Journal of Contemporary Health Law and Policy*, 10, 47–68.
- Pols, J. (2015). Towards an empirical ethics in care: Relations with technologies in health care. *Medicine, Health Care and Philosophy*, 18(1), 81–90. <https://doi.org/10.1007/s11019-014-9582-9>
- Region Hovedstaden. (n.d.). *Fordele og ulemper ved brystkræftscreening*. Region Hovedstaden. Retrieved May 7, 2025, from <https://www.regionh.dk/Sundhed/Screening/Brystkr%C3%A6ftscreening/Sider/Fordele-og-ulemper.aspx>
- Region Hovedstaden. (2023). *Redegørelse om kapacitetsmæssige udfordringer på brystkræftområdet til Task Force for Patientforløb for Kræft- og Hjerterområdet*. Region Hovedstaden. [https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerter/Bilag-2023\\_39-Region-Hovedstadens-redegoerelse\\_-brystkraeftomraadet.ashx](https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerter/Bilag-2023_39-Region-Hovedstadens-redegoerelse_-brystkraeftomraadet.ashx)
- Region Midtjylland. (n.d.). *Screening for brystkræft*. Sundhed - Region Midtjylland. Retrieved December 13, 2024, from <https://www.sundhed.rm.dk/sundhedstilbud/kraftscreening/screening-for-brystkraeft/>
- Region Midtjylland. (2023). *Redegørelse for udfordringer på brystkræftområdet*. Region Midtjylland. [https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerter/Bilag/Bilag-2023\\_42-Region-Midtjyllands-redegoerelse\\_-brystkraeftomraadet.ashx](https://www.sst.dk/-/media/Viden/Kraeft/Task-Force-Patientforloeb-Kraeft-og-Hjerter/Bilag/Bilag-2023_42-Region-Midtjyllands-redegoerelse_-brystkraeftomraadet.ashx)
- Robb, K. A., Simon, A. E., Miles, A., & Wardle, J. (2014). Public perceptions of cancer: A qualitative study of the balance of positive and negative beliefs. *BMJ Open*, 4(7), e005434. <https://doi.org/10.1136/bmjopen-2014-005434>
- Schaarup, J. F. R., Aggarwal, R., Dalsgaard, E.-M., Norman, K., Dollerup, O. L., Ashrafian, H., Witte, D. R., Sandbæk, A., & Hulman, A. (2023). Perception of artificial intelligence-based solutions in healthcare among people with and without diabetes: A cross-sectional survey from the health in Central Denmark cohort. *Diabetes Epidemiology and Management*, 9, 100114. <https://doi.org/10.1016/j.deman.2022.100114>
- Sheir, S., Manzini, A., Smith, H., & Ives, J. (2024). Adaptable robots, ethics, and trust: A qualitative and philosophical exploration of the individual experience of trustworthy AI. *AI & SOCIETY*. <https://doi.org/10.1007/s00146-024-01938-8>
- Sundhedsstrukturkommissionen. (2024). *Sundhedsstruktur kommissionens rapport – Beslutningsgrundlag for et mere lige, sammenhængende og bæredygtigt sundhedsvæsen*. Indenrigs- og Sundhedsministeriet. <https://www.ism.dk/Media/638545635292256419/Hovedrapport-tilg%C3%A6ngelig-fil.pdf>
- The Boston Consulting Group. (2017). *Analyse af kapacitetsanvendelsen på kræftområdet*. <https://www.ism.dk/Media/8/5/170407-kapacitetsanv-paa-kraeftomraadet.pdf>
- Van Cauwenberge, D., Van Biesen, W., Decruyenaere, J., Leune, T., & Sterckx, S. (2022). “Many roads lead to Rome and the Artificial Intelligence only shows me one road”: An

interview study on physician attitudes regarding the implementation of computerised clinical decision support systems. *BMC Medical Ethics*, 23(1), 50. <https://doi.org/10.1186/s12910-022-00787-8>

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

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

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

Originalartikel (uden for tema)

# Homes, Homeliness, and Otherness in Medical Anthropology “at Home”

Camilla Brændstrup Laursen<sup>1</sup>

---

<sup>1</sup>Department of Health Science and Technology, Aalborg University  
cbla@hst.aau.dk

Laursen, Camilla Brændstrup. 2026. Homes, Homeliness, and Otherness in Medical Anthropology “at Home”. *Tidsskrift for Forskning i Sygdom og Samfund*, nr. 44. 148-169. DOI: 10.7146/TFSS.V25I44.160333

Indsendt 10/25, accepteret 06/26, udgivet 06/26

*Since the 1990s, “medical anthropology at home” has been used to describe ethnographic research on health, illness, and care conducted in the anthropologist’s own society. While the notion has mainly been debated as a methodological issue within anthropology, it is perceived in this article as an analytical concern with broader relevance for discussions in qualitative health research. Drawing on ethnographic fieldwork conducted in Denmark, the article analyzes two empirical cases: one from research among people diagnosed with Irritable Bowel Syndrome (IBS) and one from research in the acute healthcare system. Focusing on concrete fieldwork situations, the article explores how experiences of being “at home”, “not at home”, and “not not at home” are negotiated, shift, and intertwine in fieldwork situations that initially appear familiar. The article draws on Hartmut Rosa’s phenomenologically inspired conceptualization of “home” (2019) and Bernhard Waldenfels’ phenomenology of the alien (2011) as analytical resources for examining how familiarity, resonance, and otherness shape ethnographic knowledge production. Through this combination of empirical analysis and conceptual reflection, the article argues that “at home” can be understood as a situational, changeable, and relational condition that shapes what can be noticed, articulated, and known in ethnographic and qualitative studies*

of health, illness, and care. By proposing a relational perspective that acknowledges that distinctions between “home” and “not home”, “familiar” and “unfamiliar” are not always clear-cut, the article invites readers to reflect on researcher involvement and the epistemic conditions of knowledge production in familiar health research settings.

## Hjem, Hjemlighed og Fremmedhed i Medicinsk Antropologi “Hjemme”

Siden 1990'erne har begrebet medicinsk antropologi “hjemme” (“at home”) været anvendt til at beskrive etnografisk forskning i sundhed, sygdom og behandling, der udføres i antropologens eget samfund. Selvom begrebet primært er blevet diskuteret som et metodisk anliggende inden for antropologien, anskues det i denne artikel som et analytisk spørgsmål med bredere relevans for diskussioner inden for kvalitativ sundhedsforskning. Med afsæt i etnografisk feltarbejde udført i Danmark analyserer artiklen to empiriske cases: én fra forskning blandt mennesker diagnosticeret med irriterende tyktarm (IBS) og én fra forskning i det akutte sundhedsvæsen. Med fokus på konkrete feltarbejdssituationer undersøger artiklen, hvordan erfaringer af at være “hjemme”, “ikke hjemme” og “ikke-ikke hjemme” forhandles, skifter og flettes sammen i feltarbejdssituationer, der umiddelbart fremstår velkendte for forskeren. Artiklen trækker på Hartmut Rosas fænomenologisk inspirerede forståelse af “hjem” (2019) og Bernhard Waldenfels’ fænomenologi om fremmedhed (das Fremde) (2011) som analytiske ressourcer til at undersøge, hvordan fortrolighed, resonans og fremmedhed former etnografisk vidensproduktion. Gennem denne kombination af empirisk analyse og begrebslig refleksion argumenterer artiklen for, at det at være “hjemme” kan forstås som en situeret, omskiftelig og relationel position, der former, hvad der kan iagttages, artikuleres og erkendes i etnografisk og kvalitativ forskning om sundhed, sygdom og omsorg. Ved at foreslå et relationelt perspektiv, som anerkender, at der ikke altid kan skelnes skarpt mellem “hjemme” og “ikke hjemme”, og hvad der er “velkendt” og “ukendt”, inviterer artiklen læseren til at reflektere over forskerens involvering og de epistemiske betingelser for vidensproduktion i sundhedskontekster, som forskeren kender personligt eller professionelt.

## Introduction

Around the 1990s, the concept of medical anthropology “at home” was discussed as a common denominator for research conducted in the anthropologist’s own (often Western) society. The concept offered a common ground for discussing the

particularities of research conducted in settings familiar to the anthropologist. These discussions highlighted both advantages of familiarity (such as ease of access and shared language) and the challenges it posed, including the risk of overlooking patterns and taking norms, practices, and assumptions for granted (Hadolt, 1998; Van Dongen & Fainzang, 1998; Van Ginkel, 1998). However, the concept also gave rise to critical questions and debates. For example, Hadolt criticized the concept of “at home” for reproducing problematic dualisms between “us” and “them”, the “West” and the “Rest” and for suppressing “non-Western” anthropological voices, the interconnectedness and the relationships between ‘we’ and ‘others’, other forms of differences than cultural ones, and similarities which cut across the ‘we-other’ and ‘home-abroad’ division” (Hadolt, 1998, p. 321). Like Hadolt, Gullestad argued in favor of overcoming these dualisms. Reflecting on the contribution of Scandinavian anthropology, she wrote:

*If anthropology is truly to become a comparative study of society and culture, modern Europe and the United States must become an integral part of the subject matter. It is necessary to overcome the now often inevitable opposition between “us” and “them,” between anthropology “at home” and “abroad” (Gullestad, 1989, p. 71).*

Reflecting on her many years of experience conducting research in her native country, Norway, Gullestad later added that anthropologists working “at home” are “seldom just insiders” (Lien & Melhuus, 2011, p. 140), and that making sharp distinctions between anthropology “at home” and “abroad” comes with the risk of missing the many connections between people living in different parts of the world. Indeed, in our globalizing world, we may ask ourselves what constitutes the difference between anthropology “at home” and “abroad” when both “home” and “abroad” are *glocal* settings. In this world, health issues continue to develop that cut across the “local” and the “global” (e.g., the COVID-19 pandemic), and healthcare systems often face similar challenges across geographical borders. These challenges include persistent shortages of health professionals, rising demands related to ageing populations and multimorbidity, and structural challenges such as overcrowded emergency departments and fragmented care pathways.

In this article, I revisit discussions of medical anthropology “at home”, exploring what “at home” may mean when ethnographic research is increasingly conducted in *glocal*, personally and/or professionally familiar settings. In doing this, I seek to shift the discussion from questions of location and insider status towards an analytical examination of “home” as a situational and relational condition that shapes knowledge production about health, illness, and care. While anthropology’s

history as a study of cultural otherness (Leistle, 2016) has rendered questions of research “abroad” and “at home” particularly salient within the discipline – sometimes normatively privileging the former as “more anthropological” than the latter (Anderson, 2021; Logan et al., 2023) – methodological and analytical considerations about being “at home” should not only interest anthropologists. As ethnographic approaches are taken up across disciplines – and as health professionals increasingly engage in qualitative research in settings they know well – reflections on the conditions under which knowledge is produced become highly important. This is especially so when attention is directed towards “patient perspectives” and “patient involvement” in efforts to address global health issues and structural healthcare challenges such as the ones mentioned above. While proximity and familiarity may ease access to study patient experiences, they do not in themselves facilitate unmediated insight into how experiences of illness, embodiment, and care are lived and articulated. By analyzing “at home” as a condition that is experienced and actively produced by researchers and interlocutors alike, this article shifts attention from accessing patient perspectives to examining the epistemic conditions under which such perspectives are produced.

In the following, I review discussions on medical anthropology “at home” before introducing my theoretical inspirations for rethinking the notion of “at home”. These are Rosa’s phenomenologically inspired, relational conceptualization of “home” (2019) and Waldenfels’ phenomenological reflections on the entangled relationships between own and alien, homeliness and otherness (2011). I use these theoretical inspirations for thinking through ethnographic examples from conducting fieldwork in and near my hometown in Denmark. Specifically, I draw on ethnographic examples from my PhD project on everyday experiences of gut trouble (Laursen, 2023) and from my postdoc project on experiences of continuity and care coordination in the acute healthcare system (Laursen et al., In review). I analyze being “at home” not only as an experiential or emotional condition, but as something that is negotiated and sometimes strategically mobilized by both researchers and interlocutors in concrete fieldwork encounters, with consequences for what can be known. Overall, I argue for a non-dichotomous understanding of medical anthropology “at home”, and for an understanding of “home” as referring to a resonant relationship with a segment of world (Rosa, 2019) rather than a fixed, geographical place. I point to the methodologically and analytically productive potentials of dwelling with the gaps, overlaps, and intertwinements between “home” and “not home” in ethnographic fieldwork.

## Revisiting medical anthropology “at home”

While “at home” has been discussed broadly within the discipline of anthropology, medical anthropology “at home” arguably constitutes an analytically distinct area insofar as questions and challenges relating to being “at home” tend to be articulated differently and/or become intensified in studies of illness, embodiment, and care. According to Hadolt (1998), medical anthropology “at home” differs from anthropology “at home” because it often examines biomedicine – a globally standardized system of knowledge and practice – thereby complicating what can be considered research “at home” and “abroad”. Other medical anthropologists have described how studying illness and suffering activates the researcher’s own body, morality, and biography in unique ways (Reis, 1998; Zaman, 2008; Larsen & Schwennesen, 2024). Most people have experiences with illness and healthcare, yet they may not know biomedical language and practice, and therefore, healthcare institutions may appear “familiar yet unfamiliar” (Van Ginkel, 1998). Whether in a hospital or in a sick person’s home, it is difficult not to become morally and bodily involved, like Zaman who, amid “life and death events” (2008, p. 148) in a Bangladeshi hospital ward could not help but console relatives waiting outside an operation room. In healthcare settings, a neutral presence often appears impossible (Zaman, 2008; Larsen & Schwennesen, 2024), as one’s presence will always be evaluated and negotiated by clinic staff, patients, and relatives (Wind, 2008). As remarked by Van Ginkel (1998), getting access to doing research in a medical institution is difficult. This may be one of the reasons why many fieldworks in clinical settings are conducted by “insiders” who managed to gain access (e.g., Zaman, 2008; Larsen & Schwennesen, 2024).

A substantial part of the literature on (medical) anthropology “at home” consists of reflexive accounts in which ethnographers seek to position themselves, arguing for why they were “natives” or not, and which strategies they used to overcome the challenges of being “native” (e.g., Munthali, 2001; Qamar, 2021; Zaman, 2008). Such reflections on positionality and accompanying methodological pros and cons are important. However, the following review will rather focus on two arguments that also surface in the literature on (medical) anthropology “at home”, and which are important to the question of what “at home” may mean when ethnographic research is increasingly conducted in personally and/or professionally familiar settings: Firstly, the argument that “home” is more than a geographical place, and secondly, the argument against dichotomous or dualist understandings of being “at home”.

Discussing the notion of “medical anthropology at home”, Hadolt argues that as a relational term, “‘at home’ receives its meanings in relation to its counterparts,

most prominently ‘abroad’, the ‘Rest’, everything which is not ‘home’” (1998, p. 316). In this context, he raises a series of questions about what “home” might mean – ranging from place and social relations to political and ideological connotations – without further unfolding them. In line with this questioning, other anthropologists emphasize that “home” has multiple meanings (Van Dongen & Fainzang, 1998, p. 245), and that “[at] homeness’ involves more aspects of one’s life than the geographical or cultural” (Reis, 1998, p. 307). According to Reis, being “at home” is about emotions, shared experiences, and processes of identification; it is about feeling “at home” and sharing experiences with one’s interlocutors through “emotional resonance” (1998, p. 304). Madden understands “home” as “a small gemeinschaftlich environment” (1999, p. 261), encompassing geographical, cultural, social, and emotional dimensions that appear familiar to someone. Nordquest (2007) distinguishes between being “at home” and being “in the field”, representing the difference between the two as a switch in one’s state of mind between “doing ethnography” and “regular life”. Similarly, Caronia (2018) argues that being “at home” and being “abroad” are epistemic positions between which the anthropologist makes cognitive shifts. Larsen and Schwennesen’s study (2024) in a Danish orthopedic surgical department adds that shifts between “home” and “not home” are not simply cognitive but also grounded in embodied and material conditions, and that the researcher’s geographical, social, and professional “at homeness” is a productive part of knowledge generation.

As indicated above, “home” is often framed in opposition to “not home”. Against this, Madden has argued that “difference and similarity, familiarity and unfamiliarity, should be seen as being inextricably linked, as ‘inseparable modalities of the same phenomenon’” (1999, p. 269). During his hometown fieldwork, Madden experienced the “familiar” and the “unfamiliar” as co-present. Similarly, Anderson (2021) describes “home” as an unstable category, marked by ambivalence and the simultaneous feeling of being both “at home” and “not at home”. Tsuda, writing from her position as a “semi-native” Japanese American, argues that anthropologists are best understood as “partial outsiders and partial insiders who experience various degrees of acceptance and cultural insight” (2015, 15). Likewise, Zhao (2017) shows how the same places may appear as both “home” and “field”, depending on one’s study object and interlocutors. With the expression “field-home”, Bilgen and Fábos propose to completely dissolve the distinction between “field” and “home”, conceptualizing “home” as “an assemblage of life stories, relationships and experiences we accumulate at different places, times and levels throughout our lives” (2024, p. 949).

While the above-mentioned literature highlights that “home” is more than a geographical place and that insider/outsider positions are always partial and

shifting, less attention has been paid to how homeliness is situationally produced and transformed, and to the consequences of these processes for ethnographic knowledge production. Rather than treating homeliness as a methodological problem to be managed, this article explores how moments of resonance and alienation may emerge within “at home” fieldwork, conditioning what can be known.

## Home, resonance, and the alien

In discussions about medical anthropology “at home”, notions of being “familiar” and “at home” are sometimes used interchangeably. However, as I will show in the analysis, being familiar with (in the sense of knowing based on previous experience) a place, person, practice, object, or situation does not always translate into being “at home”. In fieldwork situations, ethnographers may be intimately familiar with the research topic, setting, or the interlocutors prior to and beyond the fieldwork. During fieldwork, such “background” familiarities may come into play, but they do not in themselves produce experiences of being “at home”. When I turn to phenomenologically informed conceptions of “home” in the following, it is not to engage in philosophical discussions about what “home” is, but to develop an analytical lens through which the ethnographic category of “at home” can be rethought. I find the works of the German sociologist Hartmut Rosa (2019) and the German philosopher Bernhard Waldenfels (2011) particularly helpful because they enable an analytical shift from treating “at home” as a methodological position to examining being “at home” as a relational and situational experience that can be negotiated and shapes how knowledge is generated amid familiarity and otherness.

Contrary to common definitions of “home”, which are often tied to concepts of place or residence, as implied in the opposition to “abroad”, Rosa’s conceptualization of home is bound up with his idea of “resonance” as a mode or quality of human relationships to the world. While Rosa’s work is often positioned within critical social theory, I approach his concept of resonance as phenomenologically informed, drawing particularly on its affinities with thinkers such as Waldenfels. Rosa suggests that the concept of home refers to:

*[A] specific form of reference or relation to a segment of world that has been adaptively transformed – in the classic sense, a place where things speak to us and say something to us: the trees, the river, our house, or even the gas station, the industrial chimney stacks, and the local fast food restaurant. They speak because they*

*trigger resonances in our own biographical memory and the people with whom we are connected by a shared history* (Rosa, 2019, p. 359, italics original).

In this understanding, being “at home” is not to be in a specific, fixed geographical place. Rather, it is a way of relating to a “segment of world”; a relation that implies that something speaks to us by striking a chord with, for example, our identity or memories of past experiences. As formulated by Rosa: “It is the experience of a dichotomy between segments of world that have been adaptively transformed and those that remain foreign and indifferent which gives *home* meaning” (2019, p. 359, italics original). According to Rosa, there are four characteristics of resonance as a mode of being-in-the-world. Firstly, when we resonate with some part of the world (e.g., a place, an object, or another person), we are affected or touched by it. Secondly, we respond to this call by reaching out to the “segment of world” that affected us. This response need not be verbal; it can also be subtle bodily responses such as goosebumps or a high breathing rate. Thirdly, resonance affects adaptive transformation. When we resonate with someone or something, we are transformed (in more or less visible, more or less temporary ways). Rosa writes that “resonant experiences also significantly change inanimate objects (if only *for us*)” (Rosa, 2019, p. 35, italics original). He gives the example of a mountain that changes for a person when she has climbed it, compared to when she had only seen it on TV. Climbing a mountain, however, may also not cause any transformation, and in that case, the relationship is not resonant. Finally, resonant relationships are uncontrollable; we cannot plan or arrange resonance, and we cannot predict the result of adaptive transformation. For the understanding of “home”, this implies that there will always be something about “home” that we cannot fully grasp or control.

Rosa’s thinking takes inspiration from Waldenfels (2011), in particular his key point that humans are responsive beings who are required (and unable not to) respond to that which challenges us and calls upon us. Waldenfels distinguishes between “other” and “alien”, both of which could be considered counterparts to “home”. Whereas “other” refers to a difference, e.g., that a brick house is not the same as a wooden house, the “alien” (in German: “fremd”) is not simply different. The alien affects us; it is “something which seeks us out in our home (German *heimsuchen*) by disturbing, enticing, or terrifying us, by surpassing our expectations and eluding our grasp” (Waldenfels, 2011, p. 3, italics original). Referencing Freud’s point that the uncanny permeates the home, Waldenfels writes that the alien begins in ourselves and that we are therefore never entirely “at home” with ourselves (2011, pp. 76–77). The alien appears as something unfamiliar in the familiar. It distinguishes itself from our sphere of oneness, yet

the own and the alien are interwoven rather than opposed. In this sense, “home” and “alien” can be read as mutually constitutive.

Together, Rosa’s and Waldenfels’ perspectives frame “home” and being “at home” as relational and situational rather than stable or pre-given. When analyzing two fieldwork situations in the following, I use their perspectives as analytical lenses to examine how homeliness and unhomeliness are produced, overlap, and transform in concrete fieldwork encounters. Before I present the ethnographic examples, however, I will introduce the fieldworks and reflect on my position as an anthropologist “at home”.

## An anthropologist “at home”?

The empirical material presented in this article is drawn from ethnographic fieldwork conducted in the Central Denmark Region between 2016 and 2023. It is not the aim of this article to present each research project’s findings; they have been described in further detail elsewhere (Laursen, 2023; Laursen et al., In review).

I have often argued that I was doing something different than my parents when I chose to study anthropology. My parents work as a doctor and a nurse, and I have always been more familiar with the healthcare system than most people are when they simply encounter it as patients. When I began working with issues of health and illness as an anthropologist, it felt intuitive and homely. As part of my PhD project on the experience and everyday management of troublesome gut sensations diagnosed as Irritable Bowel Syndrome (IBS), I conducted fieldwork at two gastroenterology outpatient clinics and among 18 people diagnosed with IBS. When I first entered one of the clinics in 2016, I had an intuitive sense of how to speak and behave, and one of my gatekeepers (an experienced nurse) seemed to relax and consider me an “insider” when I told her that my parents worked in the healthcare system. However, I often still felt like a stranger: I did not wear a uniform, I did not have a work identity at the hospital, and I did not (at least in the beginning) speak the language of gastroenterology.

In 2022, I became employed as a postdoctoral researcher at the Prehospital Emergency Medical Services (EMS) in the Central Denmark Region. This position as a researcher in a healthcare organization provided unique insights into as well as access to the acute healthcare system, where I explored citizens’ experiences of acute care and factors affecting continuity of care. Although the Prehospital EMS now constituted my everyday research environment, the acute healthcare system

did not exactly feel like “home” to me. When I accompanied health professionals driving emergency vehicles, I took on the role of a ride-along (Seim, 2021), wearing a bright yellow vest with the tag “observer”. At times, I assisted the staff, e.g., by carrying equipment, and at times I moved around and observed staff, afflicted citizens, relatives, and bystanders. During some moments, I felt as if I was part of the prehospital staff. During other moments, I felt more “at home” with the citizens we met. Conducting research in emergency situations, I constantly judged what was appropriate for me to do and take part in according to the context and situation. My positionality could be characterized by what Wind has referred to as “negotiated interactive observation” (2008), in that my participation and observation in particular situations were not simply decided by myself but negotiated in interaction with the professionals and citizens present.

In her study on epilepsy care in The Netherlands, Reis (1998) describes how she had to make herself “a home” in an “epilepsy world” that was previously unknown to her. There are many ways in which an “IBS world” differs from an “emergency care world”. In this sense, the fieldworks I have just described required two different kinds of “homemaking”. Rather than being community studies, each research project was oriented towards a specific health challenge. While such projects receive funding and could be argued to be part of the *raison d’être* of today’s medical anthropology (Logan et al., 2023; Van Ginkel, 1998), they may also bring about challenges for the researcher who seeks to make herself “a home” while moving from one problem-oriented project to another.

In the following, I present two fieldwork situations: one from my PhD fieldwork, in which I visit and interview Natasha, a young woman troubled by diarrhea diagnosed as IBS, and one from my postdoc fieldwork in which I drive with a rapid response vehicle to help an unconscious woman. I have chosen to present these two situations because homeliness and unhomeliness, familiarity and unfamiliarity are particularly at stake, making them relevant for discussing experiences and negotiations of being “at home”.

## On gut trouble and what it means to be “at home”

Although the fieldwork that became the basis for my PhD dissertation began near my Danish hometown in 2016, one could also argue that the fieldwork began in Nepal in 2013, when I got a severe stomach infection in Kathmandu and went to a biomedical clinic where I became enrolled in a survey on traveler’s diarrhea and its possible long-term effects, including IBS. Although I never got a diagnosis, I did

continue to be troubled by my gut for years. In 2014, I was referred to a colonoscopy at a Danish private clinic, where the gastroenterologist said he wondered why so many young people, especially young women like me, complained about problems such as constipation, bloating, stomach pain, and diarrhea. Although my gut, weirdly enough, seemed to “reset” after I got another stomach infection in Nicaragua in 2015, the questions stayed with me: What was going on with this gut trouble, and why would it seem to affect young women in particular?

Fast forward to October 2016, when I am sitting in Natasha’s house in a rural village in Jutland, Denmark. At the time, Natasha is in her twenties; two years older than me. We have met during a consultation at a gastroenterology clinic, where she told the doctor about her stomach cramps and many visits to the toilet. Natasha’s troubles began when she was pregnant with the youngest of her three children. On the day of our first interview appointment, one of her daughters is at home sick, and as we sit and talk in the living room, the five-year-old girl cuddles up to her mother in the sofa, pulling up her duvet. Natasha tells me that she rarely visits her friends because she does not like to have to use their toilets often. “I don’t think that it’s fun to do it other places than at home. Well, it’s just a bit, you know... I prefer to do it at home”, she says, pausing to signal that she does not mean to elaborate on this statement.

*Natasha: You know, shit is just not a very charming thing to talk about [lort er da bare ikke særlig charmerende at tale om<sup>1</sup>]. That’s just how it is. You know, it’s not a topic that you think [det er da ikke sådan et emne, man tænker]... “Today, I would like to talk about that.” Anyway, people who know me and people whom I’m close to know it, of course, but people I talk to once in a while, you know, the kinds of people you meet in Føtex [a Danish supermarket] or in the kindergarten or something like that. They’re not the ones that you tell about it, I think [det er jo ikke sådan nogle, man står og fortæller det til, synes jeg ikke].*

*Camilla: No, it’s not really the first thing you say in the supermarket? (smiling)*

*Natasha: Not really (laughing).*

*Camilla: And this thing you told me about finding it difficult to visit other people, is that because... I mean, do you think that’s embarrassing, or?*

---

1. The interview was conducted in Danish. I provide some Danish quotes and italicize certain words to highlight linguistic nuances that are difficult to convey through translated quotes. This version of the interview extract was also presented in my unpublished PhD dissertation (Laursen, 2023).

Natasha: *Well, it's a bit embarrassing. But it's also just that thing about... You know, when you have an upset stomach, then it sometimes makes noise and it smells, and in that situation, it's just not very nice to be together with other people. Well, I think it's a bit, you know [Altså, jeg synes det er sådan lidt, du ved]... It's a bit embarrassing, it smells, and it makes noise. For sure, I don't think that it's particularly funny.*

The interview extract thematizes being “at home” in different ways. We are in Natasha’s home – a small, grey house – talking about episodes of diarrhea that she considers so intimate and embarrassing that she prefers to keep them within the confines of her home. Rather than simply expressing a preference for being at home, Natasha actively produces home as a protected space. Through her distinctions between people who are close to her and know her, and peripheral social relations (people encountered in the supermarket or kindergarten), she draws boundaries around who may enter specific domains of bodily exposure and conversation. By inviting me into her physical home and into a conversation about a topic that she otherwise considers private, Natasha actively transforms the interview into a “homely” space of trust and familiarity. This transformation is sustained through humor, shared norms of embarrassment, and my responses as a listener who signals recognition. Like Natasha and most other people in Denmark, I grew up with the understanding that it is preferable to defecate at home (Thomsen & Pröschild, 2019), and that gut-related matters are taboo in everyday conversation (an understanding that is not culturally universal, as demonstrated by, for example, Ecks’ study on gut relations in India (2003)). Natasha’s repeated use of the Danish expression “du ved” (“you know”) implies that she appeals to my understanding or expects that I already know what she is talking about. Furthermore, with the small words “jo”, “da”, and “man” (the latter coming close to “one” or “you” in English), she creates a generalizing distance to her own utterance, thus underplaying her own preferences and habits, and making her relation to (talking about) gut trouble appear as a matter of fact. More than indicators of shared understanding, however, words such as “jo”, “da”, and “man” can be argued to function strategically. Through these linguistic markers, Natasha calibrates the degree of familiarity between us and invites resonance.

Using the image of two tuning forks, Rosa explains his idea of resonance by emphasizing how it is produced when the vibration of one body makes another vibrate at its own frequency (Rosa, 2019, p. 165). Thinking with this image, Natasha’s story of gut trouble affected me by calling upon my own experiences, and although this made me relate to her experiences and get a sense of what she meant, I still did not know exactly how she felt because we, like two separate tuning forks, were both vibrating but did not share the same frequency. Although

Natasha and I were both young, Danish women, and we shared experiences of gut trouble, our lives were also very different. At the time, we both lived off public benefits, but while I received the Danish state educational support (SU) as a university student, she received cash benefits (kontanthjælp) as an unemployed person having unsuccessfully tried to find a job that she could manage with her many health problems. While I had never been hospitalized, she was hospitalized once a month or every second month. As she said: "I do feel that I am affected, you know, with the asthma, the lungs, the stomach, the back, and the wrists." Natasha lived in the countryside, while I lived in a city. At the time, she had three children and an ex-husband, while I was unmarried and had no children. With Merrild's words, we lived "parallel lives". Based on her comparative study of lower working class and higher middle-class people in Denmark, Merrild (2015) uses the expression to refer to the different lives that take place within the same cultural and discursive context, but which never intersect. Likewise, I felt that my life was running parallel with some of my interlocutors' lives. I often explained this to the interlocutors, sometimes strategically mobilizing my own story of gut trouble, by stating that I recognized their feelings, but that I was curious to know what it meant to them and why they felt that way. This helped me in my attempt to balance the merits of mutual understanding with efforts to continue exploring their lived experiences without simply appearing as "a foolish person asking silly questions" (Munthali, 2001, p. 114); a common challenge for ethnographers working "at home". Acknowledging recognition while explicitly suspending claims to full understanding formed part of an ongoing negotiation of homeliness in my fieldwork.

In another article, Meinert, Grøn, and I have analyzed the lived, embodied experience of IBS as being characterized by shifts between the gut appearing as "me", "not-me", and "not-not-me" (2022). Our analysis was inspired by Willerslev's analysis of hunting practices in Siberia (2004). Willerslev explains that the Yukaghir hunters keep a "double perspective": They strategically assume the animals' point of view while maintaining their own perspectives as hunters. By making themselves move, smell, and look in certain ways, they mimic the animals, yet they never forget their intention to kill them. In this way, they are not animals, but at the same time, they are not *not* animals. In line with this, my fieldwork could be characterized by a double perspective of being *not not* "at home". When I interviewed Natasha, I sometimes felt and acted like being "at home" (e.g., I knew the language, the places she spoke about, and I could relate to her everyday experiences of gut trouble), and at other times "not at home" (e.g., I had never been to her house before and lived quite a different life). When I visited Natasha, I felt

that I was in the “field”. Then, gradually, the further I drove away from her house and towards my own apartment, the more I felt that I left the “field” to go “home”. When I was physically “at home”, I did not perceive myself to be on fieldwork, but nevertheless, I was not *not* on fieldwork, either, as I would sometimes make field observations during my “regular life”. Nordquest (2007) has referred to this peculiar situation as a process of “switching hats”, arguing that “going on fieldwork” can be understood as a shift in one’s mindset. With inspiration from Willerslev (2004), I would rather say that I was wearing different hats at the same time, balancing between states of homeliness and unhomeliness that were always implicated in each other. Turning to the next empirical example, I will elaborate on this state of in-betweenness and how the unfamiliar may appear within the familiar.

## On an unconscious woman and the intertwining of familiar and unfamiliar

While conducting fieldwork as a postdoc employed at the Prehospital EMS, I sometimes joined emergency vehicles that were driving around in my hometown. One afternoon in May 2023, I was sitting in the back seat of a physician-staffed rapid response vehicle. I had never been inside such a car before or even interacted with one. The sirens were on, and it was difficult for me to hear where we were going, to whom, and what had happened. However, I grasped some of the short messages that the medical doctor and paramedic read aloud from the screen between them: “A woman in a blue car”, “allergic reaction”, “unconscious”, “known heart disease”, “does not respond to pain stimulation”. These were messages from the Emergency Medical Coordination Center, which had received the call to 1-1-2 (the national emergency helpline) and dispatched an ambulance and the rapid response vehicle in which we were sitting. I looked out of the window and suddenly recognized the suburban neighborhood. We were close to the primary school that I attended, and many of my classmates used to live in this area. When I went to play with them after school, we would stroll the sidewalks with our colorful school bags. While chatting about everything and nothing in particular, we would pass one brick house after another; one well-trimmed garden after the other. Sometimes, we would play a game of not stepping on the lines between the sidewalk tiles. I remember my time in the neighborhood as enjoyable; the streets felt nice and quiet, and there was an atmosphere of comfortable slowness.

On this sunny afternoon in 2023, the quiet atmosphere is contrasted by the flashing blue lights and sirens of the rapid response vehicle and the ambulance that arrived just before us. The two men from the ambulance lift a middle-aged woman out of a blue car. Her body hangs limply down, and white foam appears from her mouth. Her face is as grey as the sidewalk that they carefully place her on. The doctor from the rapid response vehicle kneels, shakes the woman by her shoulders, and shouts her name. The woman does not move, and her eyes are closed. Two women stand beside her hugging each other. They have tears in their eyes and stare anxiously at the unconscious woman whom they seem to know. One of the women has a stethoscope around her neck. She explains that the woman suddenly got sick when they were driving home from a flower shop. She talked about a burning sensation and wondered whether she could have touched some poisonous plant. Then, she suddenly passed out, and the woman driving the car quickly called 1-1-2.

While the professionals unpack medical equipment and prepare a stretcher, a few people pass by, keeping distance while directing curious gazes towards the two bright yellow emergency vehicles. I imagine where those people may be going. Some young boys look as if they are on their way home from school. A young man with a child in a cargo bike drives the other way. Perhaps he has picked up his son from the after-school center where two of my best friends used to go. I walked this stretch of sidewalk with one of those friends last summer. It was a nostalgic, but not exactly cozy walk. Her father had just died unexpectedly, much too early. It was a shock for everyone. As we walked past all the places in the neighborhood where we grew up and got to know each other, we talked about our memories of him.

I look at the health professionals, who are measuring the unconscious woman's blood pressure. It appears to be very low. They pull her shirt to the sides and attach wires to her chest. The wires are connected to a machine that measures her heart rhythms. The professionals work together to get the woman up on the stretcher and into the ambulance. They ask me to carry some equipment, and I consider where to place myself. Since I am with the professionals, I want to be close to them and the patient rather than the two women who stand around three meters away from the ambulance. At the same time, I do not want to be in the way in the ambulance in this critical situation. I end up standing around one meter from the open ambulance door. The paramedic explains to me that they are giving the woman a dose of adrenalin. I can see how she slowly opens her eyes. She looks confused and gradually begins to answer the questions posed to her by the doctor. She says that she feels nauseous and dizzy. Three of the health

professionals are in the ambulance, examining the sick woman. A lot of things are going on at the same time. They check her heart rate, measure her blood pressure, take notes to the medical record, and press down her stomach to check for internal bleeding. I look towards the two women, who are still standing on the sidewalk a couple of meters away, talking about the situation. One of them says: "It's lucky that we are so close to a hospital and that the ambulance arrived so quickly." The other woman agrees. A third woman has appeared from a nearby garden, asking what has happened. I cannot help telling them that the woman in the ambulance has opened her eyes. I know that it is not my responsibility, but it feels like an obvious thing to do.

This example illustrates how fieldworks conducted in healthcare settings particularly activate the researcher's own body, morality, and biography (Reis, 1998; Zaman, 2008; Larsen & Schwennesen, 2024). By telling the women that the patient had opened her eyes, I acted on my moral involvement in the situation. I momentarily stepped into an intermediary role, drawing on my sense of being "at home" with both the professionals and the bystanders. In the situation, I felt as if different worlds met: The stretch of sidewalk that I associated with my childhood, and the prehospital world of ambulances and rapid response vehicles in which my everyday work life unfolded. In Rosa's terms, this stretch of sidewalk had been "adaptively transformed" (Rosa, 2019) to appear as a homely place for me. I felt connected to it, and it made me recall memories of the place. However, on this day in May 2023, it transformed into a scene of accident (in Danish: "skadessted"), as formulated in the medical record. Centering on the unconscious woman, the sidewalk was suddenly a (temporary) caring place where health professionals worked, and people around the unconscious woman did their best to care for her and each other. It no longer appeared as an ordinary stretch of sidewalk that people would simply pass, but as a place for medical equipment, work plans, hugs, and worries. In this situation, my partial at-homeness in the prehospital organization functioned as a resource that allowed me to move quite freely, to be asked to carry equipment, to receive informal explanations from the professionals, and to take notes about their procedures and collaboration that I would never have noticed if I had simply passed the scene in my everyday life. Although both the street and the rapid response vehicle were familiar to me, the place appeared unhomely and strange. As a "home", the street was not a place that belonged to me or that I could control. It felt as if it was taking multiple shapes at the same time. I was not *not* home, but I was not completely "at home" either. In line with Waldenfels' thinking (2011), the unfamiliar arose within that which was otherwise familiar, and the familiar emerged in the middle of the unfamiliar. In other words, strangeness and

unfamiliarity did not emerge from me keeping a distance to what I was studying; it emerged from within the homely and familiar.

There was something uncanny about the experience of conducting this fieldwork “at home”. Driving with an emergency vehicle, I had no impact on where in the city I would go, and while we were on our way, I could not always see where we were going. This meant that I would often feel thrown into a place or situation. I might suddenly find myself in a place that appeared foreign to me, to help a person I had never seen before, but I might as well encounter my own grandmother. At one point during fieldwork, this uncanniness and collapse of the “field” and my “home” was epitomized in a dream that woke me up: I dreamt that I was in a huge venue, waiting for some show to begin when I suddenly heard someone shouting for help. People were asking around for someone familiar with resuscitation. In the dream, I rushed to help, feeling proud that I knew about emergency care and had recently taken a mandatory first aid course. I kneeled beside a person who was lying on the floor, and began the routines I had learned, pushing down rhythmically on the chest while humming Queen’s “Another One Bites the Dust” inside my head. I had been doing it for a while when someone asked if I realized who the person was. Only at that point, I saw that it was my father.

## Concluding discussion

In this article, I have revisited the concept of medical anthropology “at home”. Drawing on ethnographic fieldwork conducted in personally/professionally familiar settings, I have argued for understanding being “at home” as a situational and relational condition that shapes ethnographic attention and knowledge production about health, illness, and care.

Empirically, I have drawn on two fieldwork examples: one from my PhD research among people diagnosed with IBS, and one from my postdoctoral research in the acute healthcare system. These fieldwork situations are local, yet entangled with global contemporary healthcare. Conditions such as IBS are defined through globally circulating biomedical categories, diagnostic criteria, and treatment guidelines, while acute care systems are structured through organizational models that travel across national contexts. In this way, doing medical anthropology “at home” does not mean studying bounded local worlds, but engaging with glocal configurations. The two fieldwork examples illustrate that at-homeness is not a given, even when research is conducted in settings that initially appear familiar. In the case of Natasha, shared language, norms, and everyday references facilitated

resonance and trust, enabling intimate conversations about gut trouble. At the same time, these shared points of reference did not erase difference, but co-existed with divergent life trajectories, embodied experiences, and social positions. The prehospital case, by contrast, exemplifies how familiar places such as a childhood neighborhood may suddenly appear transformed and unhomey when it becomes a site of emergency care. Across both cases, being “at home” appeared not as a stable position but as characterized by shifts between feeling and acting like being “at home”, “not at home”, and “not *not* at home”.

Thinking with Rosa’s phenomenologically inspired conceptualization of “home” as a resonant relationship to a segment of world (2019), I have suggested that being “at home” emerges through processes of adaptive transformation: A place may feel homely in one situation and unhomey in another, or even both simultaneously. Waldenfels’ phenomenology of the alien (2011) further illuminates how experiences of otherness may arise from within the familiar. As exemplified in the prehospital case, homeliness and otherness are not opposing states; they may co-exist and intertwine. Taken together, my analyses support a relational, non-dichotomous understanding of medical anthropology “at home” that foregrounds the conditions that shape what can be noticed, articulated, and known in ethnographic studies of health, illness, and care.

## Towards a relational understanding of medical anthropology “at home”

At the turn of the millennium, anthropologists debated the concept of medical anthropology “at home” in light of broader concerns about the discipline’s future. Some questioned whether research conducted in the anthropologist’s own society would replace radical encounters with alterity (Peirano, 1998), while others worried about “home blindness” when working as an “insider” (Van Ginkel, 1998). Today, medical anthropology “at home” has demonstrated its viability, e.g., in the face of the COVID-19 pandemic that urged studies of health and illness in socio-cultural contexts but restricted travels to distant field sites (Logan et al., 2023). Supporting calls for a more inclusive discipline in which diverse methodologies and field sites can be recognized as “anthropological enough” (Logan et al., 2023), I do not suggest discarding the notion of “at home”, nor to stop discussions of key anthropological concepts such as “distance”, “field”, or “home”. Rather, I extend earlier work that emphasizes that “home” is more than a geographical place, and that insider and outsider positions are partial and shifting (Reis, 1998; Larsen & Schwennesen,

2024; Caronia, 2018; Tsuda, 2015; Zhao, 2017; Bilgen & Fábos, 2024; Madden, 1999). I do so by proposing a relational understanding of being “at home”. By this, I mean an understanding of “at home” as a situational condition that emerges through negotiation in concrete interactions and is constituted through entangled relationships of “own” and “alien”, “familiar” and “unfamiliar”, “homeliness” and “otherness”. In this understanding, medical anthropology “at home” is not defined by contrast to what it is not (i.e., medical anthropology “abroad”), but by attention to the relational dynamics through which “home” and “not home” are continuously produced, unsettled, and reconfigured in practice.

Approached in this way, “at home” is neither a background condition that precedes fieldwork nor a stable positional attribute that grants privileged access or threatens analytical distance. Instead, it appears as a situational achievement that may emerge, falter, or transform as researchers and interlocutors engage with one another, with places, and with embodied experiences of health, illness, and care. This relational understanding allows us to move beyond binary oppositions between proximity and distance or familiarity and estrangement, and to focus instead on how these dimensions are co-produced in practice, with consequences for what can be sensed, articulated, and known. By conceptualizing home as a resonant relationship (Rosa, 2019) and foregrounding the intertwining of the own and the alien (Waldenfels, 2011), the article reframes earlier concerns about “home blindness” and loss of alterity. Rather than treating proximity as a methodological and analytical obstacle, it shows how alterity may appear not despite, but from within, fieldwork conducted “at home”.

## Implications for health research and healthcare

The analysis presented in this article has potential implications beyond anthropological debates about ethnographic knowledge production. It highlights the importance of not taking homeliness for granted in either health research or healthcare practice. As healthcare increasingly moves beyond clinical institutions and into citizens’ homes and everyday settings through home-based, mobile, and community-centered services, questions about homeliness become more salient. The current Danish political focus on “proximity care” (Carstensen et al., 2025) constitutes a response to rising healthcare demands, workforce shortages, and fragmentation that challenge healthcare systems across Europe, yet the ways in which these transnational health challenges are lived and managed differ. Situated at the intersection of global healthcare rationalities and local everyday worlds,

ethnographic studies can produce knowledge about the cultural specificities of how places are adaptively transformed to both be homes and caring places. As my empirical cases indicate, neither private homes nor public spaces automatically function as caring places, nor do they necessarily remain homely once mobilized for care.

This concluding perspective invites further inquiry into how shifts between homeliness and otherness characterize patient experiences (e.g., of bodies and everyday lives that, with the onset of illness, suddenly appear unfamiliar), and shape care relations and organizational practices. For such inquiries, this article offers a conceptual vocabulary for analyzing and reflecting upon matters of homeliness and otherness that are essential to the understanding of health, illness, and care.

## Acknowledgements

I am grateful to the healthcare institutions and not least the interlocutors who let me conduct the ethnographic fieldwork described in this article. My PhD fieldwork was funded by the Graduate School of Arts, Aarhus University. My postdoc fieldwork was funded by the Novo Nordisk Foundation [NNF21OC0070185]. The Prehospital Research Council in the Central Denmark Region funded my participation in the 2023 Medical Anthropology at Home (MAAH) conference, “At-homeless? The future(s) of medical anthropology ‘at home’”. An earlier version of the article was presented at the conference, and I thank the conference participants, especially discussant Mette Bech Risør, for their constructive and thought-provoking feedback.

## References

- Anderson, E. R. (2021). Positionality, privilege, and possibility: The ethnographer “at home” as an uncomfortable insider. *Anthropology and Humanism*, 46(2), 212–225. <https://doi.org/10.1111/anh.12326>
- Bilgen, A., & Fábos, A. H. (2024). At home in the field, in the field at home? Reflections on power and fieldwork in familiar settings. *Qualitative Research*, 24(4), 935–954. <https://doi.org/10.1177/14687941231206770>
- Caronia, L. (2018). How “at home” is an ethnographer at home? Territories of knowledge and the making of ethnographic understanding. *Journal of Organizational Ethnography*, 7(2), 114–134. <https://doi.org/10.1108/JOE-12-2017-0067>

- Carstensen, K., Christensen, A. F., Fredens, M., Risør, B. W., Pedersen, I. D., Terkildsen, M., Arnth, L., & Nielsen, C. P. (2025). Proximity healthcare in Denmark: Evolving policy definitions and reform ambitions. *European Journal of Public Health*, 35(Supplement\_4), ckaf161.1590. <https://doi.org/10.1093/eurpub/ckaf161.1590>
- Ecks, S. M. (2003). *Digesting modernity: Body, illness and medicine in Kolkata (Calcutta)* [Doctoral dissertation, London School of Economics and Political Science].
- Gullestad, M. (1989). Small facts and large issues: The anthropology of contemporary Scandinavian society. *Annual Review of Anthropology*, 18(1), 71–93. <https://doi.org/10.1146/annurev.an.18.100189.000443>
- Hadolt, B. (1998). Locating difference: A medical anthropology “at home”? *Anthropology & Medicine*, 5(3), 311–323. <https://doi.org/10.1080/13648470.1998.9964566>
- Larsen, T. S., & Schwennesen, N. (2024). Beyond the insider/outsider debate in “at-home” ethnographies: Diffractive methodology and the onto-epistemic entanglement of knowledge production. *Nursing Inquiry*, 31(1), e12611. <https://doi.org/10.1111/nin.12611>
- Laursen, C. B. (2023). *Gut trouble: Irritation, experimentation, and welfare in Denmark* [Doctoral dissertation, Aarhus University].
- Laursen, C. B., Gehrt, T. B., Huibers, L., Bossen, C., & Væggemose, U. (in review). Patient-experienced continuity in acute care pathways: An ethnographic study. *BMC Health Services Research*.
- Laursen, C. B., Meinert, L., & Grøn, L. (2022). Alien guts? Exploring lives of and with irritable bowels in Denmark. *Medical Anthropology Quarterly*, 36(1), 139–154. <https://doi.org/10.1111/maq.12669>
- Leistle, B. (2016). *Anthropology and alterity: Responding to the other*. Routledge.
- Lien, M., & Melhuus, M. (2011). Overcoming the division between anthropology «at home» and «abroad» – Marianne Gullestad intervjuet av Marianne Lien og Marit Melhuus. *Norsk Antropologisk Tidsskrift*, 22(2), 134–143. <https://doi.org/10.18261/ISSN1504-2898-2011-02-05>
- Logan, R. I., Kihlström, L., & Mehta, K. (2023). “Anthropological enough?”: Reflections on methodology, challenges of doing fieldwork “at home,” and building a more inclusive discipline. *Anthropology in Action*, 30(2), 26–34. <https://doi.org/10.3167/aia.2023.300204>
- Madden, R. (1999). Home-town anthropology. *The Australian Journal of Anthropology*, 10(3), 259–270. <https://doi.org/10.1111/j.1835-9310.1999.tb00024.x>
- Merrild, C. H. (2015). *Parallel lives: Anthropological perspectives on social differences in bodily experiences of sensations and health care seeking practices* [Doctoral dissertation, Aarhus University].
- Munthali, A. (2001). Doing fieldwork at home: Some personal experiences among the Tumbuka of Northern Malawi. *African Anthropologist*, 8(2), 114–136.
- Nordquest, M. (2007). Of hats and switches: Doing fieldwork “at home.” *North American Dialogue*, 10(1), 18–20. <https://doi.org/10.1525/nad.2007.10.1.18>
- Ørsted, A. S. (in preparation). *Bedre liv med diabetes—Et casestudie om sundhedssystemers og -professionelles muligheder og barrierer for i praksis at løse et problem med social ulighed i sundhed* [Better life with diabetes—A case study of the possibilities and obstacles of healthcare systems and healthcare professionals to practically solve a problem with social inequality within healthcare].
- Ørsted, A. S., Nissen, M. A., & Vardinghus-Nielsen, H. (2025). Healthcare professionals’ identification of, and solutions to, health problems for hard to reach patients with diabetes. *Professions and Professionalism*, 15(2). <https://doi.org/10.7577/pp.6005>

- Peirano, M. G. S. (1998). When anthropology is at home: The different contexts of a single discipline. *Annual Review of Anthropology*, 27, 105–128. <https://doi.org/10.1146/annurev.anthro.27.1.105>
- Qamar, A. H. (2021). At-home ethnography: A native researcher's fieldwork reflections. *Qualitative Research Journal*, 21(1), 51–64. <https://doi.org/10.1108/QRJ-03-2020-0019>
- Reis, R. (1998). Resonating to pain: Introspection as a tool in medical anthropology "at home." *Anthropology & Medicine*, 5(3), 295–310. <https://doi.org/10.1080/13648470.1998.9964565>
- Rosa, H. (2019). *Resonance: A sociology of our relationship to the world*. Polity Press
- Seim, J. (2021). Participant observation, observant participation, and hybrid ethnography. *Sociological Methods & Research*, 53(1), 121–152. <https://doi.org/10.1177/0049124120986209>
- Thomsen, S., & Pröschold, J. (2019, [15 November]). *Vi er en nation af "hjemmeskidere": To ud af tre danskere foretrækker at lave stort derhjemme*. Politiken.
- Tsuda, T. (2015). Is native anthropology really possible? *Anthropology Today*, 31(3), 14–17. <https://doi.org/10.1111/1467-8322.12175>
- Van Dongen, E., & Fainzang, S. (1998). Editorial: Medical anthropology at home: Creating distance. *Anthropology & Medicine*, 5(3), 245–250. <https://doi.org/10.1080/13648470.1998.9964561>
- Van Ginkel, R. (1998). The repatriation of anthropology: Some observations on endo-ethnography. *Anthropology & Medicine*, 5(3), 251–267. <https://doi.org/10.1080/13648470.1998.9964562>
- Waldenfels, B. (2011). *Phenomenology of the alien: Basic concepts*. Northwestern University Press.
- Willerslev, R. (2004). Not animal, not not-animal: Hunting, imitation and empathetic knowledge among the Siberian Yukaghirs. *Journal of the Royal Anthropological Institute*, 10(3), 629–652. <https://doi.org/10.1111/j.1467-9655.2004.00205>
- Wind, G. (2008). Negotiated interactive observation: Doing fieldwork in hospital settings. *Anthropology & Medicine*, 15(2), 79–89. <https://doi.org/10.1080/13648470802127098>
- Wind, G., & Vedsted, P. (2008). Om kronisk sygdom. *Tidsskrift for Forskning i Sygdom og Samfund*, 5(9). <https://doi.org/10.7146/tfss.v5i9.1324>
- Zahavi, D., & Overgaard, S. (2014). Fænomenologisk sociologi: Hverdagslivets subjekt. In M. H. Jacobsen & S. Kristiansen (Eds.), *Hverdagslivet: Sociologier om det upåagtede* (2nd ed., pp. 169–198). Hans Reitzels Forlag.
- Zaman, S. (2008). Native among the natives: Physician anthropologist doing hospital ethnography at home. *Journal of Contemporary Ethnography*, 37(2), 135–154. <https://doi.org/10.1177/0891241607312495>
- Zhao, Y. (2017). Doing fieldwork the Chinese way: A returning researcher's insider/outsider status in her home town. *Area*, 49(2), 185–191. <https://doi.org/10.1111/area.12314>

## Forfatterliste

### *Camilla Brændstrup Laursen*

Camilla Brændstrup Laursen obtained her PhD degree in medical anthropology from Aarhus University (2023). She is currently a postdoc at the Department of Health Science and Technology, Aalborg University. Her research interests include healthcare seeking, as well as everyday embodied experiences and ideas of health, disorder, and wellbeing.

### *Casper Knudsen*

Casper Knudsen is a postdoctoral researcher at Aalborg University. His work focuses on how digital technologies and data-intensive practices reshape everyday work in healthcare organizations. Using qualitative and ethnographic approaches, he studies how professionals coordinate, adapt, and sustain data quality and continuity as new systems, standards, and performance demands are introduced.

### *Christina Egelund Antonsen*

Christina Egelund Antonsen er sygeplejerske, MLP og ph.d. med speciale i ledelse og Capacity Building. Christina er ansat ved Aarhus Universitetshospital, hvor hun arbejder med og fokuserer på at forbedre implementering, ledelsespraksis- og organisatorisk udvikling i hospitalsvæsenet. Christina er tidligere programleder for "Sikkert Patientflow".

### *Claus Bossen*

Claus Bossen er antropolog fra Aarhus Universitet og er ansat på Informationsvidenskab - også Aarhus Universitet. Han forsker i de socio-tekniske sammenhænge ml. it-systemer, digitalisering og sundhed samt brugerinddragelse indenfor feltet participatory design. Det har bl.a. handlet om udvikling og implementering af elektroniske patientjournaler, og de forandringer i arbejdsgange og organisation, som det medfører. Nyligt har fokus været på det dataarbejde, som er fulgt i kølvandet på diverse strategier for at gøre sundhedsvæsenet data-drevet, og for nuværende er forskningen rettet mod at analysere de fremmende og hindrende faktorer for at implementere kunstig intelligens i klinisk praksis.

### *Iben Mundbjerg Gjødsbøl*

Iben Mundbjerg Gjødsbøl er lektor i sundhedstjenesteforskning på Institut for Folkesundhedsvidenskab ved Københavns Universitet. Med etnografiske metoder undersøger hun, hvordan prædiktive algoritmer og præcisionsmedicin former kliniske praksisser, organiseringen og leveringen af sundhedsydelse samt patienters eksistentielle erfaringer af sygdom og behandling.

### *Jens Christian Bjerring*

Jens Christian Bjerring er professor på Institut for Kultur og Samfund på Aarhus Universitet. Han har en ph.d. i filosofi fra Australian National University og er på nuværende tidspunkt forskningsprogramleder på afdeling for Filosofi og Idéhistorie. Jens Christian er særligt interesseret i emner i filosofisk logik, formel epistemologi og AI.

### *Maja Hojer Bruun*

Maja Hojer Bruun er ph.d. i antropologi og lektor ved Pædagogisk Antropologi, DPU, Aarhus Universitet, hvor hun leder Forskningsprogrammet Fremtidsteknologi, Kultur og Læreprocesser. Hun har i over 10 år lavet etnografiske studier af nye digitale teknologier og har redigeret "Palgrave Handbook of the Anthropology of Technology" (2022).

### *Mari Serine Kannelønning*

Mari Serine Kannelønning is affiliated with Oslo Metropolitan University. She holds a PhD in social science and the introduction of AI in healthcare services, which is currently her main research area. Kannelønning has a background in information studies, Science and Technology Studies and organisation and management studies.

### *Maria Bach Nielsen*

Maria Bach Nielsen holds an MSc in techno-anthropology and is an industrial PhD fellow at Aalborg University. She is employed by Cerebriu A/S. Her research focuses on the implementation of artificial intelligence in clinical practice, with a particular emphasis on trust, contextual grounding, and co-creation in healthcare settings.

### *Michal Frumer*

Michal Frumer er antropolog, ph.d. og forsker ved Regionshospitalet Horsens, affilieret med Interacting Minds Centre, Aarhus Universitet. Hun forsker i usikkerhed og etiske dilemmaer i diagnostisk arbejde og i relation til teknologier i sundhedssystemet. Frumer har et særligt empirisk fokus på kunstig intelligens i hospitalsvæsenet.

### *Miria Grisot*

Miria Grisot is Professor in Information Systems at the Department of Informatics at the University of Oslo. She holds a PhD in Information Systems from the same university. Grisot conducts research on the sociotechnical aspects of implementing digital technologies in healthcare, with recent research on data practices, data infrastructuring, and on the introduction of AI in healthcare.

### *Mette Nordahl Svendsen*

Mette Nordahl Svendsen er professor i medicinsk antropologi på Institut for Folkesundhedsvidenskab ved Københavns Universitet. Hendes forskning omhandler de eksistentielle, sociale og politiske dimensioner af nye medicinske teknologier i sundhedsvæsenet.

### *Sune Holm*

Sune Holm er Ph.d. i filosofi fra University of St. Andrews og lektor i filosofi ved Institut for Fødevare- og Ressourceøkonomi (IFRO) ved Københavns Universitet. Sune Holms forskning fokuserer på etiske og erkendelsesteoretiske emner med særligt fokus på udvikling og implementering af medicinsk AI. Han har publiceret forskning om en bred vifte af emner, herunder algoritmisk fairness, forklarbarhed og AI-understøttet klinisk beslutningstagning. Sune deltager i flere større tværfaglige nationale og internationale forskningsprojekter om udvikling af AI til sundhedsvæsenet.

### *Tom Børsen*

Tom Børsen is a member of the research group Responsible Technology Futures (Re-TECH) at Aalborg University. In 2025, he received the SEFI Francesco Maffioli Award of Excellence for Developing Learning and

Teaching in Engineering Education for his innovative integration of ethical, social, and cultural perspectives into engineering curricula. His research focuses on translating ethical, social, and cultural elements of technological development into higher education and professional development, equipping students and practitioners with practical tools for responsible innovation. In the context of AI in health, his work is particularly relevant for analyzing how accountability, transparency, and public trust can be embedded in clinical and organizational practices.

### *Victor Vadmand Jensen*

Victor Vadmand Jensen er ph.d.-stipendiat på Institut for Klinisk Medicin på Aarhus Universitet og Regionshospitalet Silkeborg samt cand.it i IT-design og applikationsudvikling og BSc i Teknoantropologi. Victor forsker i de etiske aspekter af kunstig intelligens med fokus på, hvordan brugere i sundhedsvæsenet opfatter og handler på disse etiske aspekter.

# Retningslinjer for forfattere

Tidsskrift for Forskning i Sygdom og Samfund publicerer primært artikler, der indgår i temanumre. Vi bringer dog også artikler 'uden for temanummer', hvis de behandler relevante problemstillinger (jævnfør tidsskriftets formål).

Vi publicerer følgende typer af bidrag:

- Originalartikler
- Oversigtsartikler
- Essays
- Boganmeldelser

## Originalartikler

Vi publicerer originalartikler, der systematisk og undersøgende fremstiller egen empirisk forskning gennem en præsentation af baggrund, metode, teoretisk tilgang, analyse, diskussion og konklusion. Disse punkter er indholdsvejledende og skal ikke nødvendigvis afspejles i afsnit/overskrifter i en artikel. Der lægges vægt på:

- bidragets originalitet og betydning
- inddragelse af antropologi/medicin, humaniora/sundhedsvidenskab
- inddragelse af relevant eksisterende litteratur
- sikker sammenhæng mellem metode, analyse og fortolkning, herunder overbevisende argumentation for/diskussion af forskerens valg og konklusioner.

Omfang: Max 9000 ord inkl. dansk og engelsk abstract (max 250 ord hver), noter, referenceliste

## **Oversigtsartikler**

Vi publicerer oversigtsartikler, der med baggrund i et præcist defineret emne formidler et fyldestgørende overblik baseret på en kritisk gennemgang af eksisterende viden samt peger på fornøden ny forskning. Der lægges vægt på:

- bidragets velafgrænsede fokus og betydning
- inddragelse af antropologi/medicin, humaniora/sundhedsvidenskab
- en systematisk, kritisk inddragelse af eksisterende viden
- klar faglig positionering af forfatter
- begrundelse af valg for oversigtsvinkel (historisk, metodisk, begrebsafklarende, teoretisk, m.v.)

Omfang: Max 12000 ord inkl. dansk og engelsk abstract (max 250 ord hver), noter, referenceliste

## **Essays**

Vi publicerer essays, der formidler et emne, hvor forfatterens faglige indsigt har givet anledning til at give udtryk for en ræsonnerende holdning til et emne. Der lægges vægt på:

- at teksten er afrundet og relevant
- betydningen for ny tænkning og nye ideer
- fascinerer og udfordrer læseren med overraskende formuleringer, sammenstillinger og synspunkter

Omfang: 3000 ord inkl. noter og referencer. Abstract er ikke nødvendigt til essays.

## **Anmeldelser**

Herudover modtager vi gerne relevante anmeldelser af bøger mv. Disse må i omfang ikke overstige 1500 ord.

# Udformning af et manuskript

## Sprog

Manuskriptet skal følge gængs retskrivning på et af de skandinaviske sprog. Vi antager undtagelsesvist manuskripter på engelsk, og kun hvis der foreligger en rimelig grund hertil (f.eks. ph.d.-artikel, der skal skrives på engelsk; engelsksproget forfatter). Ved engelsk manuskript påhviler professionel sprogvidering og korrekturlæsning forfatter. Redaktionen har desuden ikke kapacitet til at korrekturlæse norsk eller svensk, og det er derfor også forfatters ansvar, at der er læst korrektur på norske eller svenske manuskripter. Vi vægter et klart, præcist sprog.

## Opsætning

Manuskripter skal indsendes som word filer og sættes op med skrifttypen Times New Roman str. 12 og halvanden linjeafstand. Kortere citater (under 3 linjer) anbringes i teksten i "citationstegn". Er de længere (mere end 3 linjer) anbringes de som selvstændige afsnit indrykket fra både venstre- og højrekant.

## Indhold i indsendelser

Manuskriptet indsendes elektronisk via tidsskriftets website. På websitet opretter man sig selv som 'author' og følger dernæst anvisningerne for submission.

Artiklerne skal være udformet efter denne standard og indsendes i følgende dokumenter:

1. Dokument: forside med forfatter (e), titel (både på originalsprog og på engelsk) og korresponderende forfatter
2. Dokument: artikeltekst med dansk og engelsk abstract, brødtekst forsynet med mellemrubrikker, citater, litteraturhenvisninger, referenceliste inkl. doi-numre samt eventuelle noter (slutnoter) og eventuelle illustrationer, tabeller og figurer med forklarende tekst og kildeangivelse.
3. Dokument: Kort forfatterbeskrivelse, som angiver uddannelse, ansættelse og forskningsområde (max 50 ord)

Herudover indtastes følgende informationer under punktet 'metadata':

1. Titel
2. Forfattere
3. Abstract
4. Referenceliste inkl. doi-numre

## Referencer

Angivelser af referencer følger principperne for APA outputstyle. Det er forfatterens ansvar, at referencelisten er korrekt opstillet. På alle referencer som har et DOI skal DOI-nummer anføres og referencelisten/litteraturlisten skal både sættes ind i 'metadata' ved indsendelse af manus samt i dokumentet med artikelteksten.

Ved angivelse af sidetal i forbindelse med citationer markeres sidetal med komma (Efternavn 1998, 207) eller med s' (Efternavn, 2000, s. 207). Valget mellem de to skal være konsistent igennem hele artiklen.

Se eksempler herunder:

Bog: Bourdieu, P. (1977). *Outline of a theory of practice*. Cambridge, UK: Cambridge University Press.

Bogkapitel: Macintyre, S., & Ellaway, A. (2000). *Ecological approaches: Rediscovering the role of the physical and social environment*. In: Berkman, L.F. & Kawachi, I. (eds.), *Social Epidemiology* (pp. 332-348). New York: Oxford University Press.

Artikel: Cheek, J. (2004). *At the Margins? Discourse analysis and qualitative research*. *Qualitative Health Research*, 14, 1140–1150. doi: 10.1177/1049732304266820.

Web-referencer skal angives med fuld URL (web-adresse) samt angivelse af hvornår web-stedet blev besøgt. Derudover angives også oplysninger om forfatter, når dette er muligt. Web-referencer bør om muligt indsættes på samme måde som øvrige referencer.

Noter anbringes efter teksten og notenummeret anbringes i teksten i løftet peltitskrift. Indskriv notereferencerne manuelt (undgå tekstbehandlingsprogrammets fod- og slutnotefunktioner).

## **Proces**

Tidsskriftet påtager sig den endelige redigering, korrekturlæsning og sætning. Kommasætningen vil blive rettet efter sprognævnets anbefalinger.

Tidsskriftet forpligter sig ikke på forhånd til at trykke artikler, hverken opfordrede eller uopfordrede. Egnede manuskripter bliver sendt i peer review. Efter peer review og revidering af manuskript kan redaktionen vurdere 1) at manus er klar til korrektur og trykning, eller 2) at manus går til en yderligere revisionsrunde på baggrund af redaktionens kommentarer.