Originalartikel

Before Deployment

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

Maria Bach Nielsen[1, 2]
Tom Børsen[1]
Casper Knudsen[3]

[1]Responsible Technology Futures, Department of Sustainability & Planning, Aalborg University
[2] Cerebriu, Copenhagen
[3]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 samarbejdsmøder, 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.

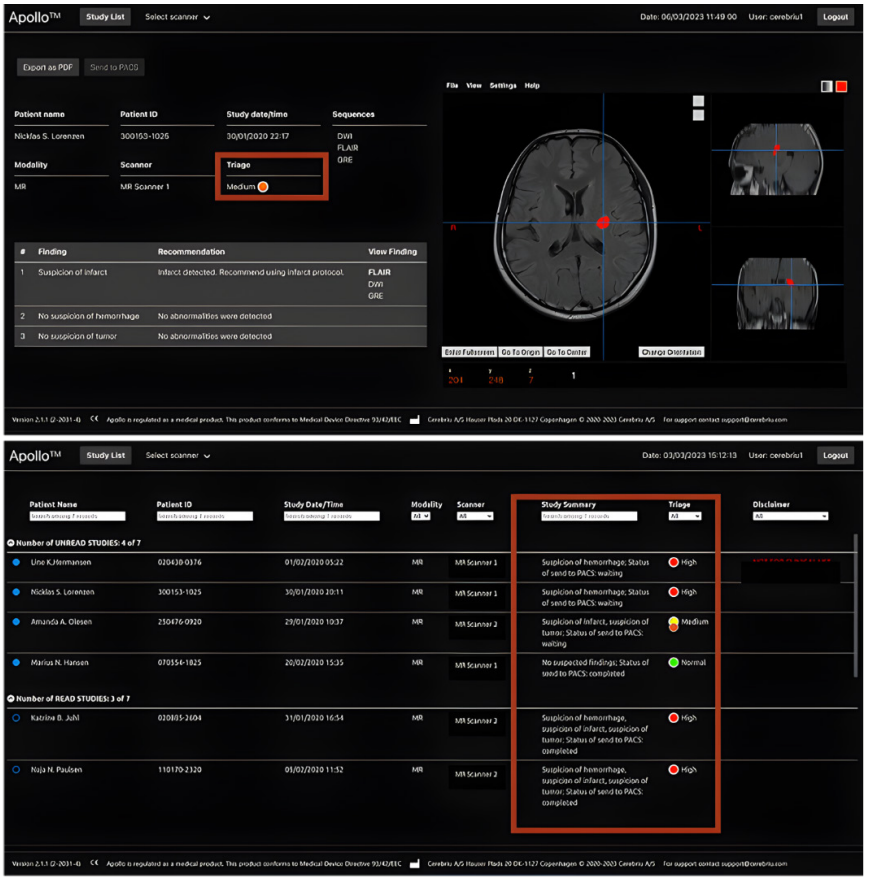

Figure 1: *Highlighted View of Apollo Version 2.1.2 Triage and Smart Backlog. Patient data shown is simulated for demonstration purposes*[1]

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*[2]

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)*

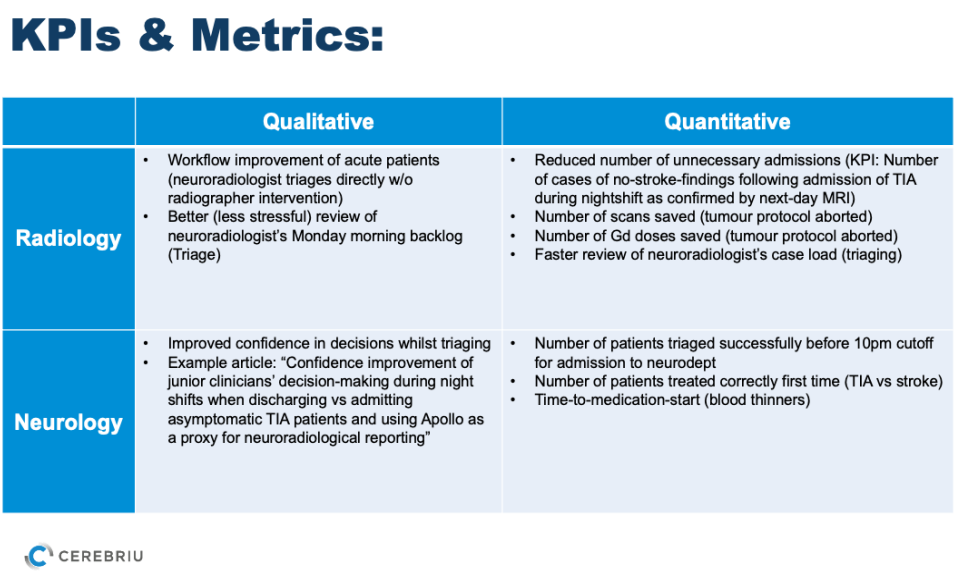

Another practical achievement of this early scoping work was the production of representations that could travel across actors and settings. Workflow maps, pilot plans, and slide decks were not neutral documentation; they were working objects that helped align what clinicians meant by workflow, what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

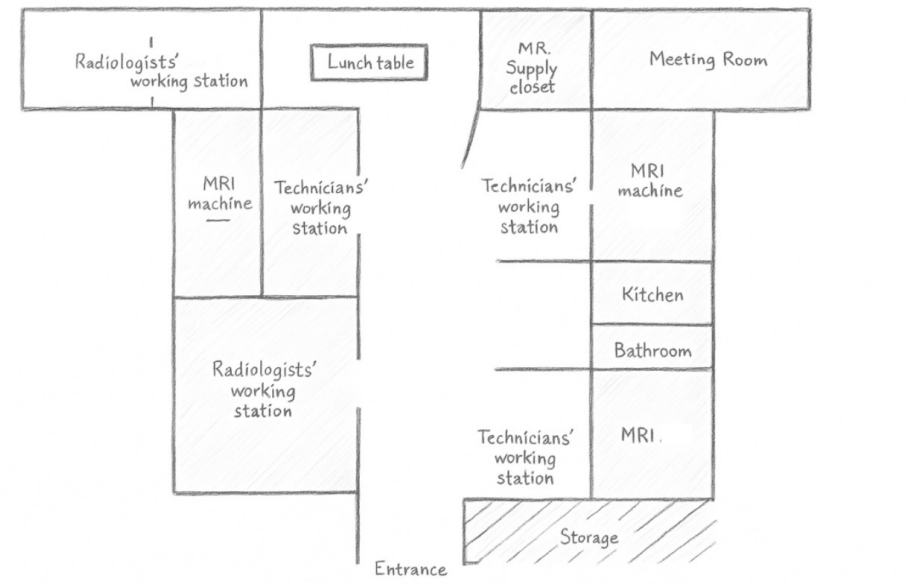

Figure 5: *Sketch of radiology department*

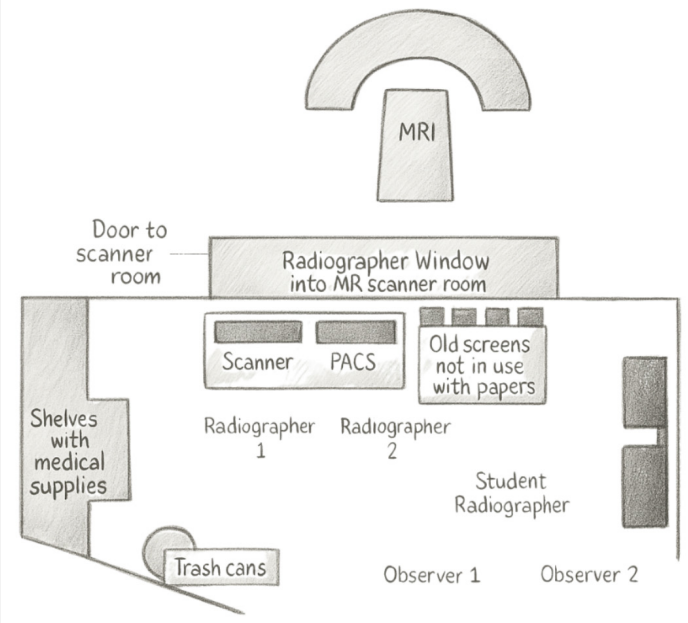

Figure 6: *The Radiographer Workstation*

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).

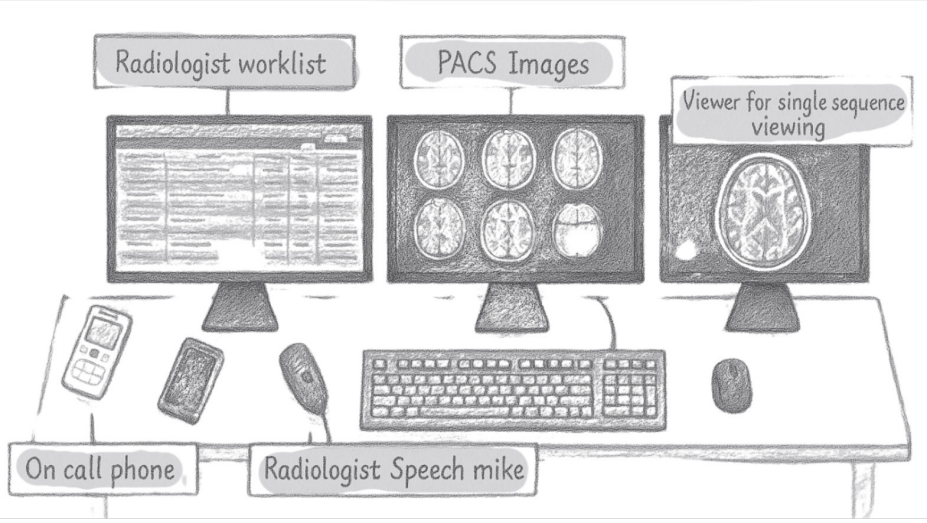

***Figure 7**: Radiologist working station*

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.

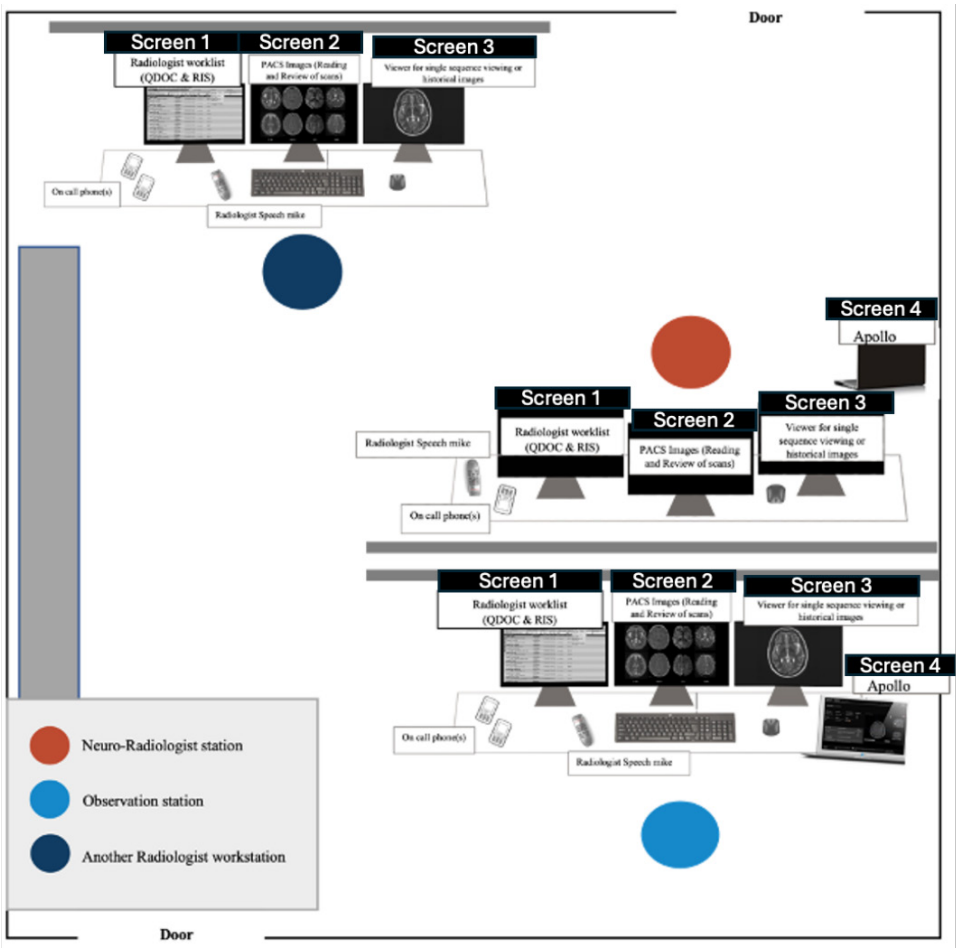

Figure 8: *Pilot observation setup (Cerebriu A/S)*

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*

Screen 1:	Patient lists and administrative systems (such as PACS and RIS, used for storing images and managing reports)
Screen 2:	The main screen for reading MRI images
Screen 3:	Comparison images or prior patient studies
Screen 4	(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., Brejnebøl, 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

McCradden, 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., McCradden, 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