

Nordic AI Meet 2026

O3. Human-Centred AI in Health & Welfare (Wed 13:00-14:00)

2. Deceptively Simple: Rethinking Task Selection for Agentic AI in Healthcare

Authors: Alina Krogstad¹

Affiliations: 1. PhD-program for Applied Information Technology, Høyskolen Kristiania, Norway

Abstract: Agentic AI are autonomous systems that identify problems and propose solutions with minimal human oversight. They are increasingly proposed as a remedy for a more efficient healthcare service but often introduce new tasks or more stress on the employees. This PhD study starts from the premise that agentic AI should only be introduced where employees themselves want it, targeting unwanted burden without adding new work or another system to manage. Over the last year I have done fieldwork at a clinic in a large Norwegian hospital, combining participatory design with a practice-based approach to articulation work, the invisible coordination labour that organizes clinical care. The data stems from 8 semi-structured interviews, over 200 hours of observation, and a participatory workshop. From this material, activity coding, staff classifying their own work into standardized categories for reporting, emerged as a compelling case. It is a task employees flagged as burdensome, and one that appears repetitive and rule-based, therefore easily assumed well-suited for agentic delegation.

My findings complicate that assumption. Staff routinely encounter categories that do not fit their actual work: codes are missing, duplicated, or indistinguishable. With no error-correction mechanism, guessing becomes normalized practice. Echoing CSCW and HCI observations that clinicians are adept at tinkering, fitting and workarounds to make ill-suited systems function, my participants develop their own personal systems and routines for coding, making the coding task characterized by tacit knowledge, individual competence and personal priorities. This is surrounded by informal social ties and varying expectations across employee groups, unevenly distributing the burden.

These findings complicate the standard automation narrative. An agentic AI trained on historical coding data would inherit years of workaround-driven, fabricated data. An agent meant to replace tacit expertise has no documentation to learn from. Rather than concluding the task resists automation, I argue my findings specify the conditions for agentic support. It must treat historical data with suspicion rather than as ground truth and deliberately elicit and preserve tacit expertise before it is lost. It must also be designed in such a way that it takes into account and safeguards existing social relationships and mediate disagreement between workers.

I contribute an empirically grounded account of why tasks that appear simple and rule-based can be poor candidates for agentic delegation, and argue human-centered AI design in healthcare must start with a participatory approach that surfaces the sociotechnical context and invisible expertise surrounding a task, before automating around it.

Keywords: human-centered AI, agentic AI, articulation work, workarounds, participatory design, healthcare