AI Comes for Nursing

Executive Summary

Artificial intelligence is reshaping healthcare as we know it. Hospital systems and healthcare technology firms are speeding up their AI adoption, despite little independent evidence that these tools improve patient care, workplace conditions, or the bottom lines of facilities. This transformation has been rapid, and, given the nature of the medical profession, the consequences can have life-or-death consequences for patients.

In this research project about the front lines of AI deployment, we examine the everyday work of nurses. Utilizing technical analysis alongside focus groups and workshops with unionized registered nurses in six states, our research maps how hospitals now use AI across nearly every part of nursing work. AI tools are used to schedule staff, admit and discharge patients, draft clinical notes, and track patients' movements. We find that AI tools are being adopted haphazardly and often without nurses' consent, with serious risks for patient care and workers' well-being. AI tools may offer cover for lean staffing models, shift responsibility onto frontline workers, and create new risks to patient safety. Moreover, despite being heralded as solutions, AI tools may do little to address the structural problems that beset healthcare systems.

Key Takeaways

- **Healthcare is a safety-critical sector, and all AI deployed in it should be treated as high risk**. There is a long-standing debate over whether policy should define high-risk AI by sector or by use case. Our research argues that healthcare applications, whether clinical or administrative, require high-risk regulatory treatment as a safety-critical sector because of the institutional context in which they are deployed. In healthcare, errors can cause immediate harm, such as patient death or rapid deterioration, as well as lasting damage to a patient's health and well-being. The healthcare sector's existing safety and accountability regimes should be strengthened to cover AI, not loosened to accommodate it.
- **Nurses carry the downstream costs of AI failures, but are shut out of the upstream decisions about procurement and deployment.** Keeping a human in the loop is increasingly positioned as a catch-all solution to the risks of integrating AI into our most sensitive systems. Our research adds to mounting evidence that human operators often get the worst of both worlds: They have little control over how systems are deployed, yet they are left to shoulder the responsibility—and even legal liability—when those systems fail. Meaningful human oversight requires independent evaluation, worker input into procurement and deployment decisions, and the ability to decide that some AI tools do not belong in healthcare at all.
- **Healthcare AI is concentrating control over patient data in dominant healthcare tech companies and Big Tech entrants, a shift that warrants greater scrutiny from competition regulators.** In the AI market more broadly, a handful of Big Tech firms control the compute, data, and market reach that the field depends on. This report adds to growing evidence that the same dynamic is taking hold in healthcare, where control over health data is cementing data monopolies and conferring enormous advantages on the companies able to exercise outsize market power to acquire patient data. Microsoft, Google, and other tech giants are acquiring clinical AI companies and partnering with large hospital systems. Meanwhile, dominant healthcare technology companies are leveraging their control over massive amounts of patient data to build AI systems, relying on health records that patients never knowingly agreed to share for that purpose. Regulators should treat healthcare data practices as a competition issue and a privacy issue at the same time, not as separate problems.