

Ethical AI for Workforce Support in Smart Hospitals

A Systems Framework Using an AI-Enabled Nutrition Support System as a Hypothetical Case Study

Kerri Reino, MBA, MS, MA, BS, AAS, FACHE, CPHQ, CPXP, PMP, SHRM-SCP

Department Administrator, Emergency Medicine & Immediate Care, Stony Brook Medicine

Short Bio

Kerri Reino is Department Administrator for Emergency Medicine and Immediate Care at Stony Brook Medicine and a PhD student in Systems Science and Industrial Engineering at SUNY Binghamton Watson School of Engineering. Her research focuses on human factors, responsible health AI, patient safety, workforce well-being, and health equity.

Abstract

Artificial intelligence is moving beyond clinical decision support into hospital operations, logistics, and workforce support. Robots already transport medications, specimens, supplies, linens, and meals. This project examines a central ethical question: Should AI compensate for organizational failures, or support healthcare professionals while those failures are corrected?

Using an AI-enabled nutrition support system as a hypothetical case study, this work examines how embodied AI could affect a hospital as a complex adaptive system. The proposed system would forecast demand, prepare individualized meals, verify allergy and dietary constraints, and route food to clinical units when staff cannot easily leave patient care. The research question is not simply whether a robot can cook, but how introducing AI changes relationships among workload, meal access, fatigue, burnout, trust, patient safety, and organizational resilience.

The ethical analysis is grounded in beneficence, nonmaleficence, justice, respect for persons, transparency, and accountability. Clinicians may miss meals because of high patient census, understaffing, inadequate break coverage, inefficient workflows, and cultures that normalize working through breaks. AI may reduce this friction, but it cannot create protected break time, improve staffing ratios, redesign workflows, or repair workplace culture. If technology makes these conditions easier to tolerate, it may treat a visible symptom while leaving the root cause intact.

The methodology combines causal-loop thinking with a pre-implementation Ethical Failure Mode and Effects Analysis (FMEA) and a governance model for calibrated trust. Preliminary FMEA scoring identifies risks extending beyond food preparation. The highest initial Risk Priority Numbers are unclear accountability after an error (320), inequitable prioritization by unit or shift (294), automation dependence and loss of manual fallback capacity (280), allergy or dietary misclassification (240), and cybersecurity or privacy exposure (224). These preliminary scores illustrate that an AI system can fail ethically, socially, and organizationally even when its technical functions appear successful.

Verification must therefore extend beyond accuracy. Proposed safeguards include human-verified allergy logic, food-safety monitoring, transparent prioritization rules, equity audits, cybersecurity testing, audit logs, human override capability, downtime drills, manual meal-access backups, incident review, and explicit stop-use criteria. Trust is treated as calibrated reliance: confidence earned through continuous verification rather than assumed because a system performs efficiently.

The significance extends beyond a robotic nutrition system. As smart hospitals delegate more essential support functions to AI, organizations must ask not only whether the technology works, but what problem it is actually solving and what responsibilities it may displace. The greatest ethical risk may not be that AI fails, but that it succeeds just well enough to reduce pressure to correct the conditions that made it necessary. Responsible AI should augment caregivers, preserve human judgment, and strengthen patient safety while reinforcing, not replacing, the institution's responsibility to care for its workforce.