SIIM CAIMI25
Conference on Artificial Intelligence in Medical Imaging
OCT 20-21 | UNIVERSITY OF CALIFORNIA, IRVINE

Automating Medical Scheduling through Agentic Assistant

Paulo Kuriki, MD, UT Southwestern Medical Center; Michael Groot, MD; Marco Pinho, MD
Ronald Peshock, MD

Introduction/Background

Coordinating clinical schedules for over 150 radiology faculty members in our institution, across multiple health systems, poses a significant logistical challenge. Each faculty member has unique clinical skills, academic commitments, and availability preferences, and coverage spans dozens of hospitals and outpatient sites. Manual scheduling requires substantial administrative effort and is prone to inconsistencies and inefficiencies.

Methods/Intervention

To address this complexity, we developed an agentic scheduling pipeline that automates the generation, validation, and revision of clinical schedules. Our system integrates with QGenda via API to download existing schedules and uses a locally deployed version of n8n as a low-code orchestration platform. The core reasoning engine is a locally hosted Mistral 3.1 small 24B model, enabling private, cost-effective LLM inference. Administrative rules are stored as natural language documents, allowing non-technical staff to define and update scheduling logic without coding skills requirement. The agentic system parses these rules, identifies inconsistencies, and proposes compliant alternative schedules for review. It is designed to scale across heterogeneous clinical environments while respecting faculty-specific constraints.

Results/Outcome

The system successfully automated major components of the scheduling workflow, including rule checking, conflict resolution, and schedule regeneration. It reduced administrative workload and improved schedule turnaround time from days to hours in test cycles. Early deployments revealed high accuracy in rule adherence and flexibility in adapting to updates. We plan to complete the feedback loop by writing validated schedules back into QGenda via API, enabling a fully autonomous scheduling pipeline.

Conclusion

Our agentic assistant represents a scalable, secure, and maintainable approach to medical scheduling in complex academic health systems. By combining local orchestration tools, LLM-based reasoning, and human-readable rule documents, we have built a system that is both powerful and user-friendly for non-technical stakeholders.

Statement of Impact

This solution demonstrates how localized, agent-driven AI infrastructure can automate high-value administrative workflows in academic radiology institutions. It empowers administrators, reduces manual errors, and lays the foundation for intelligent, rules-aware automation across broader healthcare operations.

Keywords

Agentic AI; Clinical Scheduling; Radiology Operations; Large Language Models; Workflow Automation; Healthcare Informatics