Agentic conversation on OMOP CDM: the OMCP-A2A foundation library

Niko Möller-Grell[1], Shihao Shenzhang[1], Zhangshu Joshua Jiang[7], Vishnu V Chandrabalan[7,8], Richard Dobson[1,2,3,4,5,6]

1.King's College London, Institute of Psychiatry, Psychology and Neuroscience, Department of Biostatistics and HealthInformatics

2.Institute for Health Informatics, University College London, London, UK

3. NIHR Biomedical Research Centre, University College London Hospitals National Health Service Foundation Trust, London, UK

4. Health Data Research UK London, University College London, London, UK

5. NIHR Biomedical Research Centre, South London and Maudsley National Health Service Foundation Trust and King's College London, London, UK

6. Department of Biostatistics and Health Informatics, Institute of Psychiatry, Psychology and Neuroscience (IoPPN), King's College London, London, UK

7. Lancashire and South Cumbria Secure Data Environment.

8.Lancaster University

Background

The Observational Medical Outcomes Partnership Common Data Model (OMOP CDM) has transformed healthcare analytics through the standardisation of electronic health record (EHR) data, enabling interoperability across institutions and large-scale observational studies (1). Large Language Models (LLMs) are transforming healthcare data access by replacing specialized SQL skills with natural language interfaces (2). Nevertheless, there remain challenges in linking specialised medical AI models in collaborative frameworks to tackle intricate healthcare inquiries adequately.

(3) presented foundation agents as a paradigm shift from regular language models, featuring brain-inspired modular architectures with advanced cognition, perception, and execution abilities. These systems contain five cognitive modules as a mental state: memory, world modeling, reward processing and emotion-like systems.

The framework's emphasis on specialist domain agents working together mirrors clinical team dynamics, providing advantages through parallel processing of complex medical information and the inclusion of validation mechanisms.

Google's A2A protocol provides standardized communication between independent AI agents through vendor-neutral tools for information sharing and task management (4). However, healthcare's unique requirements like stringent regulatory compliance (MHRA, HIPAA, GDPR), patient safety assurance, and domain-specific medical expertise necessitate custom augmentations beyond generic agent frameworks.

We introduce the OMCP-A2A-foundation library, a foundation-agent-based A2A protocol extension for medical use. Our system embodies the modular cognitive architecture of (3) through domain-specific agents for medical concept extraction, natural-language-to-SQL translation, and research execution with OMOP CDM data. Agents communicate via A2A protocols and execute tasks via Model Context Protocol (MCP) servers, delivering healthcare-required safety, compliance, and reliability while enabling democratized access to OMOP CDM analytics (5).

Methods

We developed the OMCP-A2A-foundation-library (https://github.com/fastomop/omcp_a2a/tree/feature/medical-a2a-framework) as an abstract architectural foundation extending Google's A2A for medical applications; implementing the brain-inspired modular design principles outlined by (3). The framework adopts a Perception-Cognition-Execution (PCE) paradigm as its core architectural pattern, providing abstract base classes (ABCs)to be implemented by concrete *MedicalAgents*.

Agent Architecture: The framework defines a *MedicalAgent* implementing the PCE cycle: perception modules process medical queries and A2A messages, cognition components integrate learning and reasoning with a dynamic *MentalState*, and execution modules generate responses through A2A messaging and MCP queries. The *WorldModel* enables domain-specific knowledge representation, supporting OMOP CDM management and patient context maintenance.

OMOP CDM Integration: Medical data models extend pydantic base classes to represent OMOP CDM structures, including *OMOPConcept*, *OMOPVocabulary*, and *OMOPConceptRelationship* classes with full vocabulary support (SNOMED-CT, ICD-10, LOINC, RxNorm). The framework provides abstract query validators for OMOP CDM database interactions.

Multi-Agent Communication: Abstract protocol classes extend A2A messaging with medical-specific extensions: *MedicalA2AProtocol* for healthcare-compliant communication, *ComplianceProtocol* for regulatory adherence, and *EmergencyProtocol* for critical situations. Agent discovery mechanisms support specialty-based routing and compliance verification through *MedicalAgentDiscovery* and *ComplianceChecker* classes.

Safety and Compliance Framework: The validation layer implements ABCs for medical safety validation (*SafetyValidator*, *DrugInteractionValidator*, *EmergencyValidator*) and regulatory compliance (*ComplianceValidator* with HIPAA and GDPR implementations). Utility abstractions provide PHI-protected logging (*MedicalLogger*), compliance metrics collection (*MedicalMetricsCollector*), and encryption (C*ryptoManager*).

Implementation Approach: The framework serves as an abstract foundation requiring concrete implementations for specific use cases. Integration with MCP servers enables database connectivity and external tool access, while A2A SDK integration facilitates agent registration, discovery, and message routing within the broader agent ecosystem.

Results

Framework Implementation: We implemented the OMCP-A2A-foundation-library, which consists of 47 ABCs organized into five architectural layers: core agent architecture (4 classes), communication protocols (8 classes), data models (12 classes), validation systems (15 classes), and utility infrastructure (8 classes). This framework provides comprehensive abstractions for developing medical agents while ensuring a clear separation between

abstract foundations and concrete implementations allowing quick development of specialized medical agents while ensuring adherence to compliance and safety standards.
OMCP-A2A’s structural components are depicted in *Figure 1*.

PCE Architecture Validation: The core *MedicalAgent* class successfully implements the PCE paradigm by defining abstract methods for the functions *perceive()*, *learn()*, *reason()*, and *execute()*. The *MentalState* management system coordinates *goals*, *emotions*, *memory*, and a *world model* throughout the agent's lifecycle.

OMOP CDM Compatibility: The data model layer of the framework represents the core structures of the OMOP Common Data Model (CDM) with full type safety, ensured through pydantic validation (6). The implementation includes vocabulary support (SNOMED-CT, ICD-10-CM, LOINC, and RxNorm) as well as capabilities for mapping relationships. The *OMOPConceptService* showcases functionalities for concept searching, navigating hierarchies, and managing vocabularies, which are critical for processing medical queries.

A2A Protocol Extension: The medical-specific protocol abstractions enhance Google's A2A framework by implementing the *MedicalA2AProtocol*, *ComplianceProtocol*, and *EmergencyProtocol* classes. The agent discovery features specialty-based routing and compliance verification mechanisms. Additionally, message routing facilitates orchestration of medical workflows, ensuring proper priority handling and adherence to audit trail requirements.

Compliance and Safety Integration: Validation frameworks offer safety-checking capabilities through abstract validators designed for drug interactions, dosage validation, and emergency detection. Compliance validators implement abstractions for MHRA and GDPR requirements, while the utility layer provides logging to protect personal health information (PHI), encryption, and an infrastructure for collecting compliance metrics.

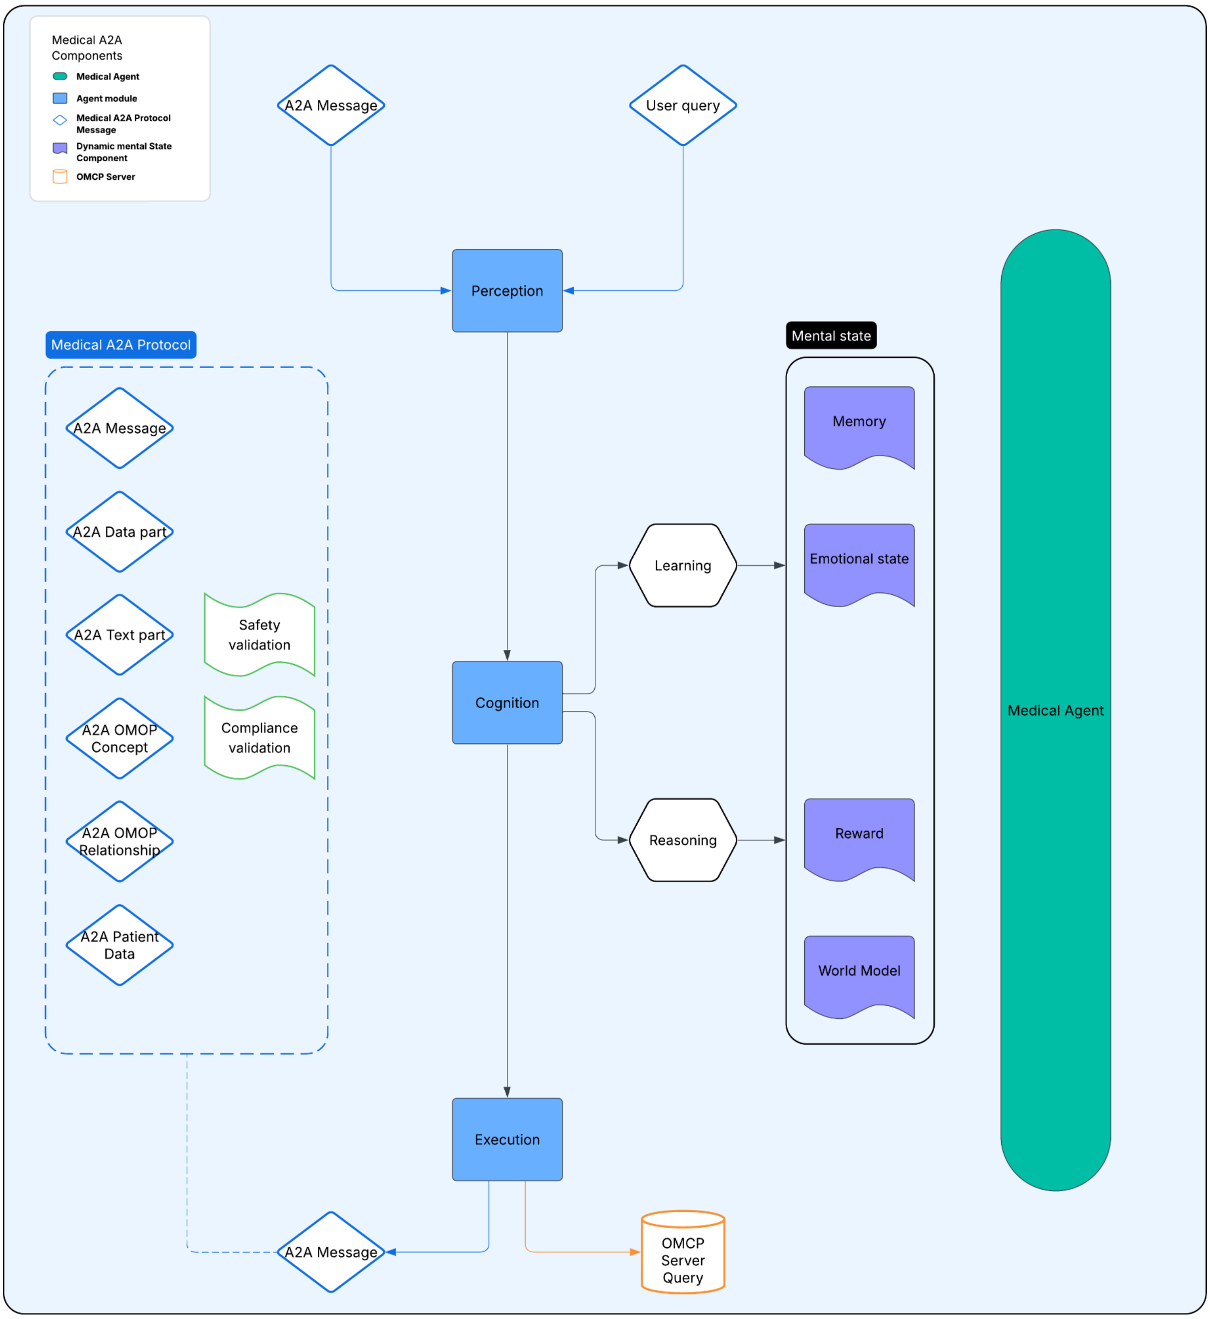

Figure 1: *Structural overview of the medical A2A conception and implementation components.*

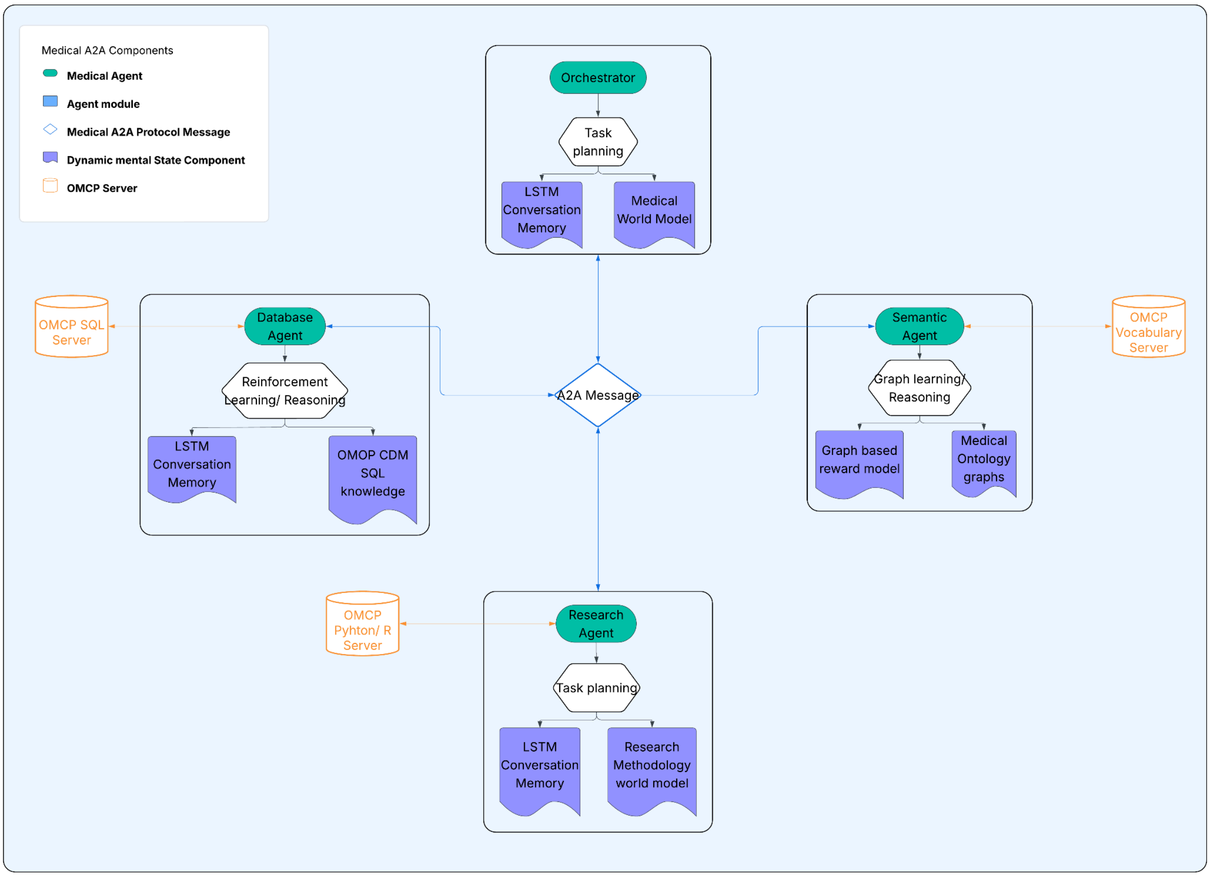

Figure 2: *Example Implementation of Medical Agent orchestration using the OMCP-A2A medical foundation library.*

Figure 2 shows an implementation of the omcp-a2a-foundation library for live medical Q&A. Three *MedicalAgents*collaborate to provide accurate live insights into OMOP CDM structured EHR data. Managed by an orchestrator agent, they collaboratively retrieve and process OMOP data to provide medical valid insights. Tool execution is handled through calls to OMCP servers while the communication between participating agents is handled through OMCP-A2A.

A POC implementation without advanced cognition using small language models (SMLs) resulted in 60% executable queries with 46.7% being medically valid, showing the potential of agentic frameworks. We will present results of the advanced implementation as outlined in *Figure 2* at the conference evaluating against sampled ACHILLES queries as well as MedHELM and SDBench benchmarks (7–9).

Conclusion

The OMCP-A2A library shows the modularisation capabilities of agentic frameworks while incorporating safety and compliance validation at protocol level. We hope to extend and refine the library's capabilities to serve as a foundation library for the implementation of safe and modularised medical agents in the future.

References

1. Observational Health Data Sciences and Informatics. The Book of OHDSI [Internet]. [cited 2025 May 4]. Available from: https://ohdsi.github.io/TheBookOfOhdsi/

2. Schuemie M, Chen Y, Fridgeirsson E, Kim C, Reps J, Suchard M, et al. Assessment of Pre-trained Observational Large Longitudinal models in OHDSI (APOLLO).

3. Liu B, Li X, Zhang J, Wang J, He T, Hong S, et al. Advances and Challenges in Foundation Agents: From Brain-Inspired Intelligence to Evolutionary, Collaborative, and Safe Systems [Internet]. arXiv; 2025 [cited 2025 Jun 2]. Available from: http://arxiv.org/abs/2504.01990

4. google/A2A [Internet]. Google; 2025 [cited 2025 May 2]. Available from: https://github.com/google/A2A

5. GitHub [Internet]. [cited 2025 May 2]. Model Context Protocol. Available from: https://github.com/modelcontextprotocol

6. pydantic·PyPI [Internet]. [cited 2025 Jul 1]. Available from: https://pypi.org/project/pydantic/

7. Bedi S, Cui H, Fuentes M, Unell A, Wornow M, Banda JM, et al. MedHELM: Holistic Evaluation of Large Language Models for Medical Tasks [Internet]. arXiv; 2025 [cited 2025 Jul 1]. Available from: http://arxiv.org/abs/2505.23802

8. Nori H, Daswani M, Kelly C, Lundberg S, Ribeiro MT, Wilson M, et al. Sequential Diagnosis with Language Models [Internet]. arXiv; 2025 [cited 2025 Jul 1]. Available from: http://arxiv.org/abs/2506.22405

9. OHDSI/Achilles [Internet]. Observational Health Data Sciences and Informatics; 2025 [cited 2025 Jul 1]. Available from: https://github.com/OHDSI/Achilles