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

MCP-Driven Multi-Agent RAG System for Prostate Cancer Temporal Summarization, Lifespan Prediction, and Treatment Recommendation

Meenakshi Amulya Jayanti, University of Chicago; Zoe Calianos, University of Chicago; Hyunji Kim, University of Chicago; Brian Li

Introduction/Background

Longitudinal prostate cancer records contain rich but fragmented data that challenge clinicians' ability to make timely, explainable decisions. Existing AI systems often hallucinate, lack temporal reasoning, and fail to integrate structured and unstructured data effectively. We present a novel multi-agent system built on the Model Context Protocol (MCP), designed to produce clinically grounded temporal summaries, lifespan predictions, and treatment recommendations from synthetic longitudinal prostate cancer records.

Methods/Intervention

Our architecture uses CrewAI agents interacting via MCP to coordinate a retrieval-augmented generation (RAG)-enhanced summarizer, supervised ML models, and hallucination-resistant validation. We generated a synthetic dataset of 500 longitudinal prostate cancer cases using LLMs and expert-in-the-loop review. Each case includes PSA levels, PI-RADS scores, weight trends, bone pain severity, and treatment history. For treatment recommendation, we trained and deployed XGBoost classifiers; for lifespan prediction, we evaluated multiple regressors, selecting Ridge Regression based on RMSE. The system is hosted via FastAPI and outputs downloadable, agent-verified reports.

Results/Outcome

Preliminary evaluation on synthetic test data shows robust integration and performance across components. The summarizer produces coherent, temporally accurate narratives and citations using a custom PubMed RAG tool. The validator agent flags hallucinations and missing fields. The XGBoost classifier achieved high accuracy, with PSA and metastasis status as top predictors. Ridge Regression for lifespan prediction achieved RMSE of 0.499, MAE of 0.370, and R^2 of 0.825. A front-end interface allows physicians to input patient IDs and receive structured PDF reports with summaries, treatment recommendations, and lifespan predictions.

Conclusion

We demonstrate a reproducible, explainable, and modular agent-based LLM system capable of synthesizing longitudinal oncology data into structured, evidence-supported clinical insights. Its architecture enables plug-and-play extensibility, real-time API deployment, and interpretability aligned with clinical reasoning patterns.

Statement of Impact

This work advances the use of LLM-based tools in oncology by delivering temporal reasoning, hallucination detection, and multimodal decision support through reproducible, clinically relevant methods. It has potential to reduce documentation burden and improve patient outcomes in real-world settings.

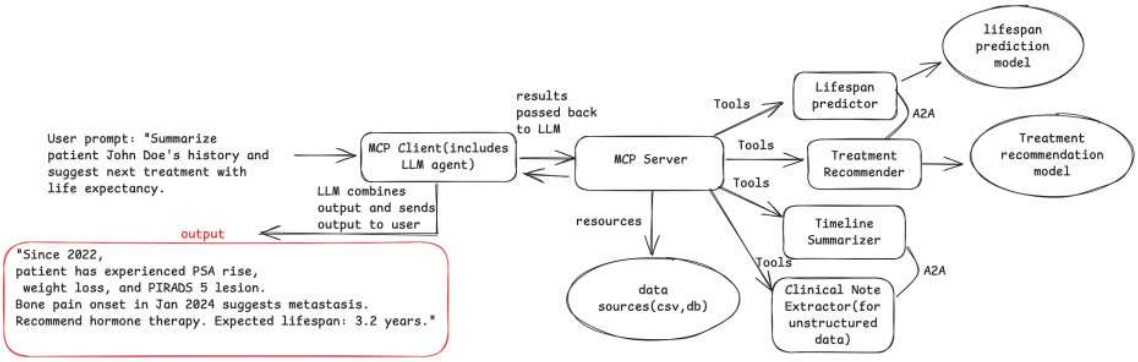

Figure 1: Model Architecture of the system

PatientID, Clinical Notes

1,"2020-03-13: DIAGNOSIS - Initial PSA=14.7 ng/mL, PI-RADS=5, Initial Weight=73.7 kg | 2020-03-13: Treatment=ADT, PI-RADS=5, PSA=14.7 ng/mL, Weight=73.7 kg, Bone Pain=Mild | 2020-06-11: Treatment=ADT, PI-RADS=5, PSA=14.5 ng/mL, Weight=74.5 kg, Bone Pain=None | 2020-09-09: Treatment=ADT, PI-RADS=5, PSA=14.4 ng/mL, Weight=75.3 kg, Bone Pain=Mild | 2020-12-08: Treatment=ADT, PI-RADS=5, PSA=14.1 ng/mL, Weight=76.2 kg, Bone Pain=Mild | 2021-03-08: Treatment=ADT, PI-RADS=5, PSA=13.9 ng/mL, Weight=76.9 kg, Bone Pain=Mild | 2021-06-06: Treatment=ADT, PI-RADS=5, PSA=13.6 ng/mL, Weight=77.7 kg, Bone Pain=Mild"

2,"2022-04-11: DIAGNOSIS - Initial PSA=15.4 ng/mL, PI-RADS=3, Initial Weight=87.8 kg | 2022-04-11: Treatment=Surgery, PI-RADS=3, PSA=15.4 ng/mL, Weight=87.8 kg, Bone Pain=None | 2022-07-10: Treatment=Surgery, PI-RADS=3, PSA=15.0 ng/mL, Weight=87.8 kg, Bone Pain=None | 2022-10-08: Treatment=Surgery, PI-RADS=3, PSA=15.0 ng/mL, Weight=87.8 kg, Bone Pain=None | 2023-01-06: Treatment=Surgery, PI-RADS=3, PSA=14.7 ng/mL, Weight=87.8 kg, Bone Pain=Mild | 2023-04-06: Treatment=Surgery, PI-RADS=3, PSA=14.5 ng/mL, Weight=87.8 kg, Bone Pain=Moderate | 2023-07-05: Treatment=Surgery, PI-RADS=3, PSA=14.4 ng/mL, Weight=87.8 kg, Bone Pain=Mild"

3,"2021-09-14: DIAGNOSIS - Initial PSA=20.0 ng/mL, PI-RADS=3, Initial Weight=91.4 kg | 2021-09-14: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=91.4 kg, Bone Pain=None | 2021-12-13: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=92.5 kg, Bone Pain=None | 2022-03-13: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=93.6 kg, Bone Pain=None | 2022-06-11: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=94.1 kg, Bone Pain=Severe | 2022-09-09: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=95.2 kg, Bone Pain=Moderate | 2022-12-08: Treatment=Surgery + Radiation + ADT, PI-RADS=3, PSA=20.0 ng/mL, Weight=96.2 kg, Bone Pain=Moderate"

Figure-2: Synthesized Data from ChatGPT

Summary for Patient P1

1. Clinical Summary for Patient 1 — Patient ID 1, diagnosed with prostate cancer, has been on Androgen Deprivation Therapy (ADT) since March 2020. Over the course of approximately one year, the patient's PSA levels have shown a gradual decrease from 14.7 ng/mL to 13.6 ng/mL. The PI-RADS score has consistently been 5, indicating a high likelihood of clinically significant prostate cancer. Concurrently, the patient's weight has increased from 73.7 kg to 77.7 kg. Bone pain was initially mild, disappeared briefly, but then returned and remained mild.

2. Clinical Literature Context — According to Peppercorn et al., 2025, the management of prostate cancer with ADT can lead to various metabolic changes, which might explain the gradual weight increase observed in the patient. The study emphasizes the importance of monitoring metabolic parameters and adjusting treatment protocols to mitigate potential adverse effects (from *Application of DNA Methylome Analysis to Patients with ME/CFS*).

3. Patient Timeline:

- **2020-03-13**: Diagnosed with prostate cancer; Initial PSA=14.7 ng/mL, PI-RADS=5, Weight=73.7 kg; Started ADT; Bone Pain=Mild.
- **2020-06-11**: On ADT: PSA=14.5 ng/mL, PI-RADS=5, Weight=74.5 kg; Bone Pain=None.
- **2020-09-09**: On ADT: PSA=14.4 ng/mL, PI-RADS=5, Weight=75.3 kg; Bone Pain=Mild.
- **2020-12-08**: On ADT: PSA=14.1 ng/mL, PI-RADS=5, Weight=76.2 kg; Bone Pain=Mild.
- **2021-03-08**: On ADT: PSA=13.9 ng/mL, PI-RADS=5, Weight=76.9 kg; Bone Pain=Mild.
- **2021-06-06**: On ADT: PSA=13.6 ng/mL, PI-RADS=5, Weight=77.7 kg; Bone Pain=Mild.

4. Validation Comments:

- All information provided in the summary and timeline is directly derived from the structured clinical data provided.
- No hallucinations, missing values, or inconsistencies were noted in the summary or timeline.
- The citation from the TrajectoryEvidenceTool was used to provide context on the potential metabolic effects of ADT, relevant to the patient's weight gain.

Treatment Recommendation:

ADT

Lifespan Prediction:

6.353760920871498

Figure-3: Sample Temporal Summary for Patient 1

MCP-Driven Multi-Agent RAG System for Prostate Cancer Temporal Summarization, Lifespan Prediction, and Treatment Recommendation - Abstract

Keywords

Temporal Summarization; Model Context Protocol; Multi-Agent System; RAG-LLM; Explainability; Clinical Decision Support, Lifespan Prediction, Treatment Recommendation