Split-Brain Auto-Planning a Hierarchical Multi-Agent LLM Framework for Inverse Optimization

Ramesh Boggula, PhD

Wayne State University, Detroit, MI

2026 JULY 19–22 VANCOUVER, BC

JOINT AAPM | COMP MEETING

ABSTRACT

Introduction: To demonstrate a novel "Split-Brain" Multi-Agent Large Language Model framework for autonomous inverse planning in radiation oncology.

Methods: We developed a dual-agent architecture within an open-source inverse planning framework (PortPy). A Reasoning Agent (DeepSeek-R1, Attending Physicist role) reviews dose volume histogram summaries and spatial dosimetric metrics to generate high-level clinical directives, such as prioritizing PTV coverage versus hotspot control. An Action Agent (Qwen-2.5, Planner role) converts these directives into explicit updates to objective function weights and penalty terms, then re-runs optimization iteratively. We evaluated the system on multiple lung cancer cases initialized with suboptimal objective weights. Performance was assessed by improvements in PTV D95%, hotspot control, and maintenance of organ-at-risk doses within tolerance limits.

Results: Across lung cases, the framework consistently improved plans starting from intentionally suboptimal weights, reaching prescription-level target coverage while strictly maintaining critical organ (e.g., Spinal Cord, Heart) constraints. The system demonstrated autonomous clinical reasoning by shifting focus to hotspot control when doses exceeded a clinical threshold (115% of prescription). In a representative case, the framework dramatically improved PTV D95% within a few iterations and subsequently reduced hotspots by autonomously escalating overdose penalties. Once clinical goals were met, plan metrics showed no further meaningful change with additional iterations, indicating stable convergence without continued retuning.

Conclusion: The Split-Brain LLM framework demonstrates the feasibility of autonomous clinical reasoning for treatment planning. By framing decisions as explicit strategies instead of opaque "black-box" behaviors, this approach provides a transparent, robust path to next-generation automated treatment planning.

CONTACT

Ramesh Boggula, PhD, DABR
Wayne State University / Karmanos Cancer Institute
4100 John R, Detroit, MI 48201
rboggula@wayne.edu
https://rameshboggula.com

INTRODUCTION

Current automated treatment planning systems often rely on static algorithms or single-model AI approaches that may struggle to balance competing clinical goals, such as target coverage and normal-tissue sparing, while providing limited transparency into planning decisions.

We introduce a novel "Split-Brain" multi-agent architecture that separates strategic reasoning from numerical optimization, mimicking the clinical workflow in which an attending physicist guides a planner. This study evaluates the feasibility of the framework for autonomous inverse planning in lung cancer cases.

METHODS AND MATERIALS

A dual-agent framework was implemented in PortPy, an open-source inverse treatment-planning platform. Lung cancer cases were initialized with theoretical objective weights and optimized using PortPy's convex optimization solver.

After each optimization, the system calculated plan-quality metrics, including PTV D95%, PTV maximum dose, spinal cord maximum dose, heart mean dose, and lung mean dose. DeepSeek-R1, acting as the Reasoning Agent or "Attending Physicist," evaluated these metrics together with the current objective weights and predefined planning goals, then generated a natural-language planning strategy.

Qwen2.5, acting as the Action Agent or "Planner," translated the strategy into numerical updates to the corresponding PortPy objective weights. The updated weights were applied, and the plan was reoptimized iteratively.

For the representative case, the workflow consisted of an initial optimization followed by three agent-guided re-optimization cycles.



Split-Brain optimization workflow. After each PortPy optimization, plan-quality metrics are evaluated by DeepSeek-R1, which generates a natural-language planning strategy. Qwen2.5 translates the strategy into numerical objective-weight updates for the next optimization cycle.

Agent roles at a glance

	Reasoning Agent	Action Agent
Model	DeepSeek-R1 [2]	Qwen2.5 [3]
Clinical role	"Attending Physicist"	"Planner"
Reads	Plan metrics vs. clinical planning goals	The natural-language directive
Produces	High-level clinical directive in natural language	Objective weight & penalty-term updates; re-runs optimization
Types	Prioritize PTV coverage; control hotspots	Increase PTV_MIN weight; escalate overdose penalty

RESULTS

The Split-Brain framework began with a target-dominant weighting strategy, assigning full priority to the PTV and no initial weighting to the OAR objectives. After each optimization, the Reasoning Agent reviewed the plan metrics and directed the Action Agent to modify the objective weights based on unmet planning goals.

In the evaluated lung cases, the Split-Brain framework improved target coverage toward the prescription level while maintaining spinal cord and heart doses within predefined planning limits. In the representative case, PTV coverage improved as the system iteratively increased the PTV underdose objective weight. When the maximum dose exceeded 115% of the prescription, the Reasoning Agent shifted priority from coverage to hotspot control, and the Action Agent increased the overdose objective weight. Once the planning goals were met, additional iterations produced no meaningful changes in the evaluated metrics, indicating stable convergence.

DISCUSSION

The Split-Brain framework used two specialized LLM agents to guide iterative inverse planning. Separating clinical strategy from numerical weight adjustment made the decision process easier to follow because each change was tied to a stated planning priority.

The study was limited to lung cases in an open-source research setting and did not include comparison with human planners or clinical treatment plans. Broader validation is needed before clinical use.

CONCLUSIONS

The Split-Brain framework guided iterative inverse planning in the evaluated lung cases. Separating reasoning from weight adjustment made the planning process more transparent.

[1] Jhanwar G, Tefagh M, Taasti VT, et al. Portpy: An Open-Source Python Package for Planning and Optimization in Radiation Therapy Including Benchmark Data and Algorithms. *AAPM 65th Annual Meeting & Exhibition*, 2023.

[2] Guo D, Yang D, Zhang H, et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*. 2025;645:633–638.

[3] Yang A, Yang B, Zhang B, et al. Qwen2.5 Technical Report. *arXiv:2412.15115*; 2024.

© Poster Template by Genigraphics® 1.800.790.4001 www.genigraphics.com