Towards LLM-based generalized automated treatment planning for locally advanced rectal cancer

Jiacheng Liu[1,2]; Kaining Yao[2]; Meijiao Wang[2]; Yichen Pu[2]; Haizhen Yue[2]; Yi Du[2]; Hao Wu[2]; Ruoxi Wang[2]; Lisheng Geng[1]
[1]School of Physics, Beihang University, Beijing 102206, China. [2]Key Laboratory of Carcinogenesis and Translational Research (Ministry of Education/Beijing), Department of Radiation Oncology, Peking University Cancer Hospital & Institute, 52 Fucheng Road, Beijing 100142, China.

2026 JULY 19–22 VANCOUVER, BC
JOINT AAPM | COMP MEETING

ABSTRACT

The growing cancer burden demands generalizable automated treatment planning (ATP). We present a large language model (LLM) - driven ATP framework for locally advanced rectal cancer that addresses the poor generalizability of data-driven approaches. Departing from data-intensive paradigms, our method uses few-shot learning of LLMs to build a prescription-agnostic ATP system. The LLM enables three core capabilities: distilling general planning knowledge from a minimal set of example plans, generating adaptive plan initializations via an LLM-based adapter module, and providing explicit step-by-step rationales for inverse planning decisions. We evaluated the framework on 45 patients with substantial variation in target volume definitions and clinical prescriptions, using independent blinded expert assessment. The ATP generated clinically acceptable plans for all patients within a median of 6 optimization iterations. Knowledge extracted from only three cases proved robust across diverse prescriptions, yielding dose distributions comparable to or superior to manual plans. The LLM's transparent reasoning trajectory enhanced explainability, and blinded experts rated automated plans as clinically superior or equivalent to manual plans in 86.7–88.9% of cases. This work establishes a generalizable, explainable foundation for reducing clinical workload and advancing automated treatment planning.

CONTACT

Jiacheng LIu
Beihang University
Beijing 102206, China
Jiacheng_liu@buaa.edu.cn

INTRODUCTION

Recent advances in modulating techniques like volumetric modulated arc therapy (VMAT) have increased treatment planning complexity, which still relies heavily on manual iteration. Automated treatment planning (ATP) offers a solution, but current data-driven methods, such as knowledge-based planning (KBP), require extensive high-quality data and lack generalizability across diverse clinical prescriptions. Large language models (LLMs) provide a novel paradigm for capturing implicit knowledge from limited data without additional training, showing promise in radiotherapy tasks. In locally advanced rectal cancer (LARC), radiotherapy regimens vary substantially—including short-course, long-course, and simultaneous integrated boost—yet existing ATP frameworks rarely accommodate such diversity, especially when lateral pelvic lymph node involvement or dose escalation is required. To address this gap, we propose a framework that leverages an LLM to achieve generalized ATP. Unlike prior studies, ours is the first to use LLMs to distill explicit planning knowledge from iterative historical data, augmenting the planning workflow. The core innovation lies in learning foundational, generalized planning logic to handle complex and highly variable clinical scenarios.

METHODS AND MATERIALS

The automated treatment planning (ATP) workflow mirrors the clinical pipeline in two stages: plan initialization and inverse planning. In initialization, a large language model (LLM)-based prescription adapter generates configuration files (machine parameters, optimization objectives, clinical requirements) tailored to diverse prescriptions. During inverse optimization, an LLM-based agent emulates an experienced radiation oncologist and physicist, evaluating plan quality and dynamically adjusting objectives based on dosimetric metrics. This agent operates under a "Learning-to-Empower" framework: in the Learning stage, an LLM knowledge extractor distills explicit planning rules from multi-iteration historical data of a few cases; in the Empowerment stage, the agent applies this knowledge to iteratively refine the plan until clinical acceptability. The knowledge extractor learns evaluation criteria and optimization strategies (direction selection, parameter type, adjustment magnitude) from prompts and iterative planning data. The prescription adapter, primed with five representative cases (covering SIB, LCRT, SCRT, LPLN), translates clinical text into planning configurations. Implementation used three sample patients for knowledge distillation (prescription: 45 Gy/50 Gy in 25 fractions) and five cases for adapter construction. The system invoked Eclipse TPS via ESAPI, exchanged data via JSON files, and terminated after 15 iterations (failure if not clinically acceptable). The LLM used was DeepSeek-R1.

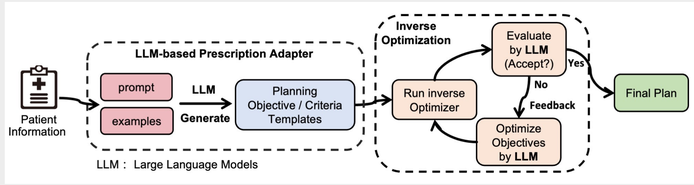

Figure1. The workflow of the LLM-based generalized automated treatment planning.

RESULTS

The proposed ATP framework successfully generated clinically acceptable plans for the validation set of 45 patients across different prescriptions, with a median of 6 optimization iterations. Total planning times were 47±16 minutes. The LLM-based knowledge extractor inferred robust planning strategies from only three cases. For the validation set across diverse prescriptions, blinded expert evaluation revealed a consistent preference for automated plans across all three reviewers. Specifically, expert 1 rated 51.1% (23/45) of cases as "automated better," 11.1% (5/45) as "manual better," and 37.8% (17/45) as equivalent. Corresponding ratings from Expert 2 were 75.6% (34/45), 13.3% (6/45), and 11.1% (5/45), and from Expert 3 were 66.7% (30/45), 13.3% (6/45), and 20.0% (9/45). Additionally, DVH comparisons across cases with varying prescriptions demonstrated the robust performance of the ATP approach, as depicted in Fig. 3. Automated plans consistently produced dose distributions comparable to or superior to those of manual plans. Notably, in the cases shown in Fig. 3(c), (e), and (f), the ATP approach achieved substantial dose reductions to the OAR, highlighting its ability to generate high-quality plans across diverse clinical scenarios.

Evaluation Criteria:

1. **Target Structures (PTV etc.)**
 - ✅ **Hard constraints met**: `AchievedDose` ≥ `HardConstraint` for *all targets*.
 - *Rationale: Target coverage is non-negotiable; iterations stop only when prescription doses are fully achieved.*
2. **OARs (FemoralHead, etc.)**
 - ✅ **Hard constraints met**: `AchievedDose` ≤ `HardConstraint`.
 - ✅ **Stable optimization**: `OFV` ≥ 0.1 for *all OARs*.
 - *Rationale: Further dose reduction is unnecessary if OAR objectives are no longer improvable.*
3. **Global Constraints (BODY, etc.)**
 - ✅ **Hard constraints met**: `AchievedDose` ≤ `HardConstraint`.
 - *Rationale: Global dose limits (e.g., body hotspots) must not exceed safety thresholds.*

Optimization Strategy:

1. **Optimization Direction**
- **Targets (PTV, etc.)** :
Adjust when `AchievedDose < HardConstraint`.
- **OARs (FemoralHead, etc.)** : Adjust when `OFV < 0.1`, indicating the current objective is easily met, allowing further dose reduction.
- **Global Constraints (BODY, IrradVolume)** : Adjust when `AchievedDose > SoftConstraint`, indicating hotspots/normal-tissue doses exceed tolerance.
2. **Parameter Adjustment**
2.1 **Adjustment Type Selection**
- **Priority Adjustments** :
Used for: - Target structures (to enforce coverage)
- Global constraints (to suppress hotspots/normal-tissue doses)
Rationale: Increases optimization weight to force dose redistribution.
- **Dose Adjustments** :
Used exclusively for OARs.
Rationale: Directly lowers the dose objective to exploit "dose room" when OFV is low.
2.2 **Adjustment Value Calculation**
Targets : AdjustmentValue = (HardConstraint - AchievedDose) × 0.5
OARs : AdjustmentValue = - (Current Objective Dose × 0.1)
BODY : AdjustmentValue = (AchievedDose - SoftConstraint) × 0.5
IrradVolume : Fixed `AdjustmentValue = 10` when triggered.

Figure2. The extracted planning knowledge with LLMs from few-shot planning samples.

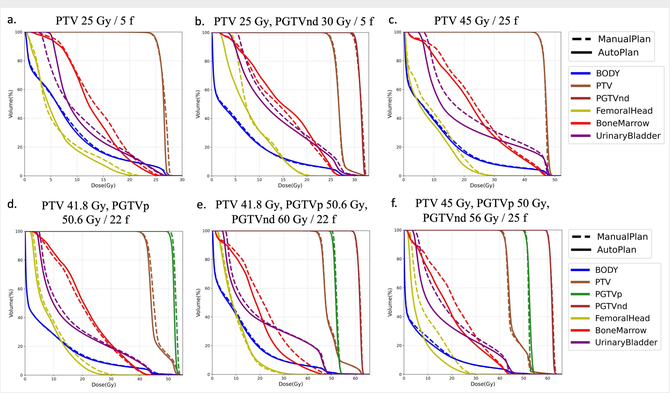

Figure3. Dose-volume histogram comparison between manual and automated plans for diverse clinical prescriptions.

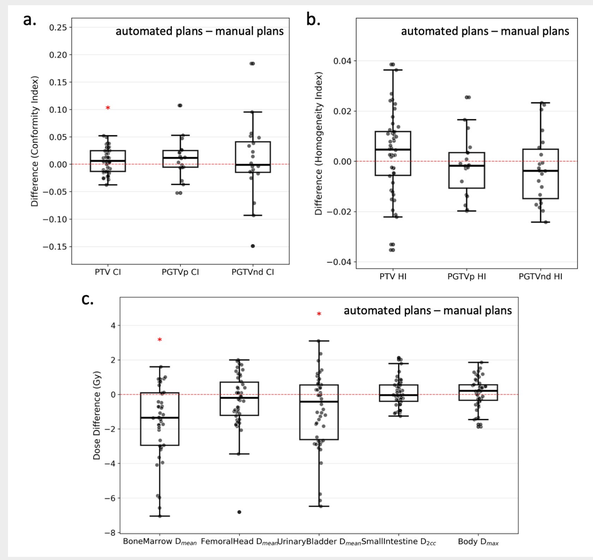

Figure4. Difference of dosimetric indicators (automated plan - manual plan) for the validation sets. Red asterisks indicate statistically significant differences.

DISCUSSION

1. Planning Knowledge Distillation
The core contribution is leveraging an LLM to extract explicit, structured planning logic from limited historical data, replicating experienced planners' reasoning rather than memorizing case specifics. The distilled rules remain human-readable, enabling direct expert review. Validation on internal and external datasets confirms the robustness and generalizability of the extracted knowledge.

2. Explainability & Reliability & Generalizability
The LLM-based ATP maintains a transparent optimization trajectory—recording planning state, assessment, and operation at each step. This mirrors the clinical manual-planning workflow, thereby fostering clinician trust and facilitating clinical adoption. A prescription adapter accommodates diverse clinical prescriptions without requiring scenario-specific retraining. No significant differences in planning time or iteration counts were observed across different scenarios, confirming the framework's broad applicability.

3. Limitations & Future Work
The current implementation depends on a cloud-based proprietary LLM, raising data privacy concerns. Future work will adopt a locally deployed LLM to ensure data security. Although the validation dataset is larger than those in prior studies, multi-institutional testing on a more extensive cohort is needed to fully assess robustness and generalization.

CONCLUSIONS

Our LLM-driven ATP framework provides a generalizable solution for generating high-quality treatment plans. It establishes a foundation for reducing clinical workload, improving planning efficiency, and advancing the development of generalized, explainable ATP systems.