Knowledge-Anchored Intelligent Multi-Agent AI Plan Check Platform to Safeguard Online Adaptive Radiotherapy

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

Ying Zhang*, Jingying Lin, Chris Kabat, Hailun Pan, Justin Visak, Fan Chi Su, Viktor Iakovenko, Siqiu Wang, Ruiqi Li, Steve Jiang, Arnold Pompos, Mu-Han Lin

MAIA Laboratory and Department of Radiation Oncology, UTSW, Dallas TX, USA

MAIA Laboratory
UT Southwestern Medical Center
Radiation Oncology

Background

Online MR-Linac adaptive radiotherapy (ART) requires rapid, complex treatment planning and QA.Small planning errors can lead to:

- Workflow interruptions
- Longer on-couch treatment times
- Increased patient safety risk

Current QA is largely manual and rule-based, making it:

- Time-consuming
- Reviewer-dependent
- Limited in detecting uncommon or previously unseen errors

Fast, comprehensive QA is especially critical for MRIgART.

Objective

- the first knowledge-anchored, multi-agent AI system for comprehensive ART plan QA.
- Support both offline planning and online adaptive workflows.
- Learn institutional best practices from curated historical plans.
- Automatically:
  - Summarize plan feature, optimization challenges, and contour list to be edited online
  - Detect planning errors
  - Flag parameter deviations
- Improve workflow efficiency, safety, and planning consistency.

CONTACT

Ying Zhang, Ph.D., DABR

Ying.zhang@UTSouthwestern.edu

METHODS AND MATERIALS

❑ Layer 1: Learning Agent (Knowledge Base&Working Memory)

- Curated institutional best practices based on online efficiency and robustness.
  - 27 approved pancreas plans
  - 64 prostate plans, selected
- AI distillated historical human plan-check comments

❑ Layer 2: RadChecker: Multi-Agent Plan Checker and Reporter

- Contour Agent, Plan agent,
- IMRT optimization agent, Dose agent,
- Critic Agent, Synthesizer Agent

❑ Web-based Adaptive Radiotherapy (ART) Board

- Displays plan information and longitudinal trends
- AI audit report:
  - plan summary for online communication
  - a list of contours to be edited online
  - AI flagged critical findings, warnings,
- physicist chat box with intuitional memory

❑ Performance Evaluation

- Plans with intentionally introduced errors.
- Quantitative and qualitative evaluation
- clarity, readability, and overall satisfaction

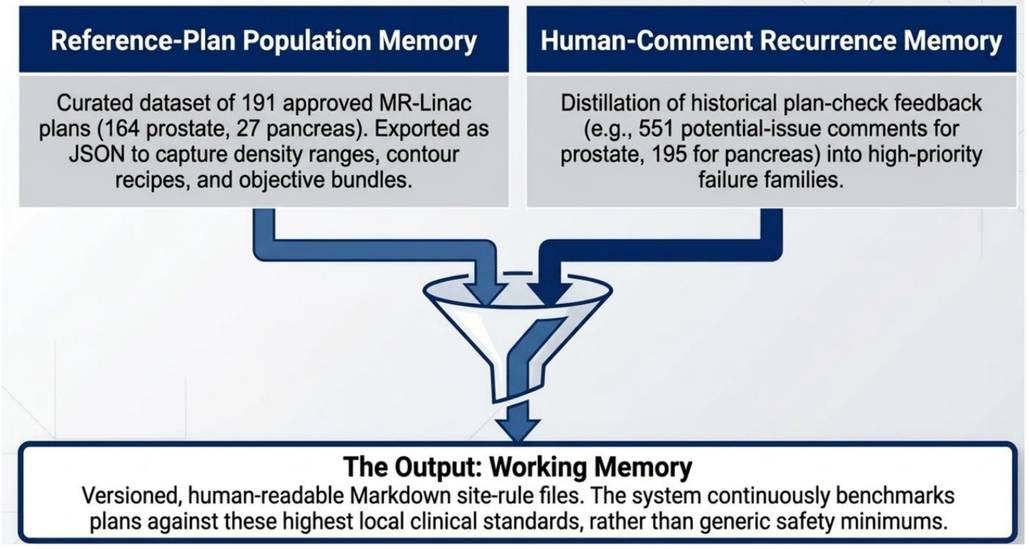

Figure 1: Learning Agent Architecture

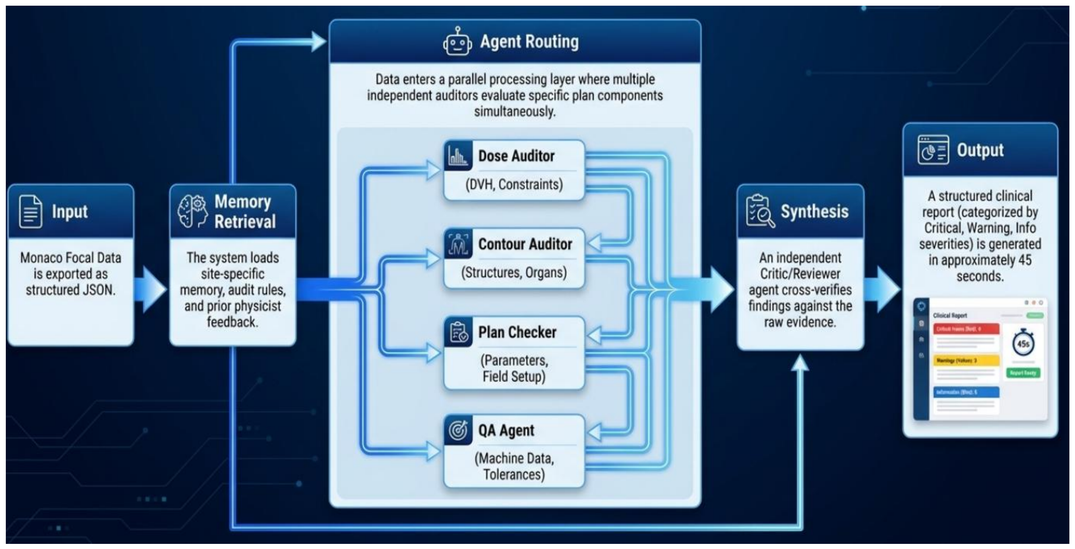

Figure 2: RadChecker System Flow: Parallel Multi-Agent Audit Process

RESULTS

❑ Two AI audit modes

- Regular Mode: ~40–50 s/plan (offline workflow)
- Fast Mode: ~20–30 s/plan (online workflow)

❑ Error detection performance

- Evaluated on 3 test plans with 46/47 errors detected
- Sensitivity: 97.9%; Precision: 93.9%
- One missed minor deviation and three false positives

❑ Qualitative evaluation

(three physicists, 5-point Likert scale)

- Report clarity: 4.3 ± 0.6
- Readability: 4.2 ± 0.3
- Overall satisfaction: 4.1 ± 0.6

CONCLUSIONS

- We developed an integrated platform combining an ART review board with an intelligent multi-agent AI auditor for online ART plan checks.
- By encoding expert human planning knowledge into the agents, the system transfers institutional expertise into an automated AI reviewer that enforces consistent institutional standards on every plan.
- The system enables fast, reliable, and intelligent plan review, addressing a critical safety and efficiency gap in offline planning and online adaptive workflows.