Redefining MRI:
From Image to Insight

ICMRI 2026
THE 14th INTERNATIONAL CONGRESS ON MRI
& 37th ANNUAL SCIENTIFIC MEETING OF KSMRM

AUGUST 27 - 28, 2026
Grand Walkerhill Seoul, Seoul, Korea
https://icmri.ksmrm.org/2026m

인공지능의료영상 (AI for Medical Imaging & Applications) 1

SY03-3

Vision-Language Model 과 Agentic AI 를 활용한 판독 보조 시스템

Songsoo Kim

Seoul National University Hospital, Korea

Vision-language clinical AI increasingly depends on probabilistic reasoning, whose outputs vary across runs and whose rationale is difficult to audit—a fundamental barrier to clinical adoption, which demands reproducibility and traceability. Longitudinal follow-up of multiple brain metastases exemplifies the challenge: each patient harbors numerous lesions that must be tracked over time, and because their responses diverge after radiosurgery, even the human report defaults to qualitative descriptions such as "partly increased, partly decreased" and fails to reach a decisive verdict. We address this by decoupling vision from agentic orchestration: a deterministic segmentation model performs perception, while a probabilistic LLM agent reasons over its outputs under the constraints of tool use, rule-based verification, and human oversight, and is thereby integrated into a deterministic clinical pipeline.

Specifically, a fine-tuned brain-metastasis segmentation model delineates and measures each lesion and tracks its trajectory across timepoints. Above it, the LLM agent layer constrains probabilistic reasoning through four deterministic mechanisms: an orchestrator that invokes tools (lesion table, tumor burden, trajectory, verdict), rule-based hard gates that resolve unambiguous cases a priori, a schema-validated structured verdict, and a human-in-the-loop that re-integrates the radiologist's interactive corrections. This layer resolves the probabilistic output into a reproducible, auditable three-class verdict (Progression / Stable / Indeterminate).

We evaluated three-class progression assessment on 311 3D brain MRI cases (up to 182 lesions and 7 timepoints per case) with a Gemini-3.5-flash backbone. Hard gates resolved 49% of the cohort deterministically prior to any agent invocation, and the full pipeline achieved 81.0% three-class accuracy, with progression-versus-rest sensitivity of 0.89 and specificity of 0.85. Residual error was concentrated in the Indeterminate class (F1 0.24), and degraded terminations—schema violations or step-limit fallbacks—accounted for only 4%. In addition, radiologists can revise and confirm the report through interactive correction.

Overall, the agent quantitatively integrates the segmentation model's measurements to deliver reproducible verdicts that clinicians are often reluctant to render. Verifiable per-lesion evidence and a human-in-the-loop grounded in interactive correction promote calibrated reliance rather than uncritical automation. This work provides an agentic design framework for integrating probabilistic LLM reasoning into a deterministic clinical pipeline.

Keywords: *Vision-language model, Agentic AI, Brain metastasis. Human-in-the-loop*