SIIM CAIMI26
Conference on Artificial Intelligence in Medical Imaging
OCT 26-27 | University of Pennsylvania, Philadelphia, PA

Smart Atlas - Search Your Archive with Visual Index

Barbara Klaudel, PhD Student, Emory University; Aleksander Obuchowski, MA

Briefly describe what attendees will see during your demo.

We will show several examples of searching for similar cases to the currently analyzed one

Problem & Motivation: What clinical or operational problem does your MVP address?

Hospitals accumulate decades of radiology knowledge: cases confirmed by biopsy, surgery, long-term follow-up, or subsequent imaging. Yet this institutional experience remains largely inaccessible during routine clinical work. The problem is no longer a lack of information. Radiologists now have access to an expanding universe of online resources, textbooks, teaching files, and digital institutional archives. Existing PACS systems can retrieve every abdominal CT performed on a 64-year-old woman in the past five years, but they cannot identify the few prior cases that resemble the lesion currently analyzed. Especially when those cases are described using different terminology or contain the answer in pathology, operative, or follow-up records. As a result, radiologists must manually search textbooks, online references, and their own memory for analogous cases. This is slow, cognitively demanding, and difficult to sustain in high-volume clinical environments. The challenge is particularly significant for trainees and early-career radiologists, who have less personal experience to draw on and often lack a timely feedback signal. The learning loop is also incomplete. Surgical findings, biopsy results, and follow-up imaging frequently fail to reach the radiologist who interpreted the original examination. Without systematic feedback, clinicians may rarely learn whether an initial interpretation was correct or how an uncertain case ultimately resolved. We do not want to automate diagnosis or remove clinical judgment. Clinicians should remain responsible for diagnosis. Instead, they should be given broader, relevant context at the moment it is needed: visually and clinically similar prior cases, their confirmed outcomes, and the evidence that connects them.

What You Built: Describe your MVP, tool, or workflow in plain language

Smart Atlas is a visual search engine for a hospital’s own radiology archive. It helps radiologists find relevant prior cases while they are interpreting a study, giving them broader clinical context without attempting to diagnose the current case. Input: The radiologist selects a full study, an image series, a single slice, or an optional region of interest within the normal reporting workflow. Smart Atlas uses the same DICOM images already stored in and served by the hospital’s PACS. What it does: a specialized medical image encoder converts the study into an embedding space where visually similar findings cluster naturally, regardless of patient or accession. Nearest-neighbor retrieval runs against a continuously updated vector index of the hospital's historical archive. Output: Smart Atlas produces a ranked list of visually similar prior cases. Each result includes the relevant images, the original radiology report, and highlighted diagnostic terms. When available, it also shows the case’s confirmed downstream outcome from the electronic health record, such as biopsy results, surgical pathology, or follow-up imaging. The radiologist can open the complete prior study in PACS with one click. The same workflow could later support clustering similar studies for batch review of one focused reading session, such as small lung nodules, and linking imaging patterns to relevant guidelines, protocols, or literature. By design, Smart Atlas does not generate an impression, flag pathology, or classify a study as normal or abnormal. It retrieves relevant evidence and institutional experience. The radiologist interprets the case and makes the diagnosis.

Tech Stack & Development Approach: What tools, frameworks, or methods did you use to build it?
Encoder: domain-specific medical image encoder producing embeddings. The encoder was fine-tuned by our team extending CLIP-like model to be sensitive to how to subtle changes in report findings correspond to visual differences, Retrieval: approximate nearest-neighbor vector index over the hospital archive, updated offline, continuously as new studies are reported; sub-second queries, QDrant vector database Data layer: read-only DICOM ingestion from PACS; outcome linkage from EHR reports and pathology Backend: Python, FastAPI, deployable on a single on-premise GPU server - 100% local processing, zero data leaves the hospital Frontend: For the demo, we used an inhouse developed web-based, zero-footprint DICOM viewer built on Cornerstone3D with smart atlas built as an extension accessible via sandboxed iframe API. Training and validation dataset: UMIE - open-source unification developed by our team of 20 public radiology corpora (882,774 images, 3 modalities) into a single RadLex-aligned annotation schema, CC-BY-NC-SA, on GitHub

Validation & Evidence: Have you done any informal testing or validation? If so, what did you find?
We have validated the core retrieval mechanism, but not yet the complete Smart Atlas workflow in clinical practice. A preliminary version of the approach was evaluated in RAD-SRAC, presented at the GenAI4Health Workshop at AAAI 2025. The study asked a controlled question: does providing a model with its nearest visually similar prior cases at inference time improve diagnostic classification? Across three imaging tasks: kidney tumor CT, chest X-ray, and brain tumor MRI, and seven vision-language models, retrieval produced reproducible improvements without model fine-tuning. The retrieval mechanism evaluated in that work is closely related to the one underlying Smart Atlas, although the study tested model performance rather than radiologist performance. We have also developed a reproducible data foundation through UMIE datasets. This allows other researchers to reconstruct the index and independently evaluate retrieval behavior. The public Smart Atlas demo runs the actual retrieval pipeline on open datasets. Informal walkthroughs suggest that search is fast and that many returned cases are radiologically plausible. Some of the most interesting results are “near misses”: cases that appear similar but differ in a clinically important feature, potentially making them useful for differential diagnosis. We have also shown the prototype to approximately 20 radiologists and collected early, informal feedback. However, we do not yet have a clinical reader study, prospective deployment, or measured evidence that Smart Atlas improves reporting accuracy, confidence, education, or time per case. Establishing real-world clinical value and identifying situations where retrieval may distract or bias readers, is the next major validation step.

Feedback You're Seeking: What specific feedback, help, or collaboration are you hoping to get from the CAIMI26 community?
The most valuable outcome would be finding a clinical champion and pilot site. We are looking for a radiology department willing to select one narrow, well-defined use case, and evaluate Smart Atlas against its own archive in an on-premise, read-only environment. We are also seeking guidance on the design of a credible clinical reader study. Key questions include the minimum number and experience level of readers, the most appropriate crossover or multireader design, suitable endpoints for accuracy and time, and how to measure possible harms such as anchoring on a retrieved case. We want to distinguish situations in which retrieval provides useful evidence from those in which it may create distraction or false reassurance. A second priority is an integration reality check. We would value collaboration with PACS, reporting-platform, or viewer vendors

that could help us understand the technical and commercial requirements for embedding visual search directly into the radiologist's existing workflow. A separate application is unlikely to achieve sustained use. We also hope to conduct in-depth interviews with potential early adopters to identify where Smart Atlas would provide the greatest and least clinical value. We are particularly interested in which subspecialties, findings, and workflow stages are best suited to similarity-based retrieval. Finally, we welcome candid feedback on clinical utility, product positioning, regulatory and validation strategy, and whether this should initially be developed as a clinical tool, an educational platform, or research infrastructure. Introductions to pilot partners, startup programs, experienced advisors, and mission-aligned early-stage investors would also be valuable.

Known Limitations & Honest Failures: What isn't working yet, or what have you already tried that failed?

The image encoder is still evolving. A close embedding does not always imply meaningful clinical similarity: cases may cluster because of anatomy, scanner characteristics, acquisition protocol, image quality, or incidental features rather than the finding of interest. We are working to make retrieval more clinically specific. We are also experimenting with "similar but importantly different" cases, which may be valuable for differential diagnosis, but we do not yet have a validated way to balance close matches with informative contrasts. To date, we have worked primarily with open-source datasets. These datasets are valuable but frequently contain incomplete labels, inconsistent metadata, limited outcome information, and occasional errors Our experiments were possible largely because of UMIE, which standardizes public radiology datasets into a shared format and ontology. Harmonization improves usability, but it cannot fully correct errors or missing information in the source data. We have not yet established performance across hospitals, countries, patient populations, scanner manufacturers, protocols, and equipment generations. Domain shift, scanner variability, and equitable representation would require broader datasets and rigorous external validation. We have previous experience integrating with HIS and PACS systems in Poland, but we have not yet completed a US integration. Effective use will require access to PACS images and reliable links to reports, pathology, surgery, and follow-up data. It may also require vendor-supported DICOMweb, HL7 or FHIR interfaces, authentication, audit logging, and viewer integration. Finally, we have no prospective evidence that Smart Atlas improves clinical accuracy or efficiency. We are discussing potential studies with several US institutions.

Potential for Impact: If this MVP were fully developed and deployed, who would benefit and how?

Radiologists could access their institution's accumulated experience while interpreting a case. Instead of manually searching multiple references, they could review visually similar studies and their confirmed outcomes within the reporting workflow. This may reduce search burden, broaden the differential diagnosis, and help readers calibrate confidence against relevant precedent. Trainees and early-career radiologists could learn from systematic, confirmed examples rather than isolated teaching cases. A resident could retrieve cases resembling the study they are currently reading, including difficult examples that initially misled more experienced radiologists. Linking the original interpretation with pathology or follow-up could restore a feedback loop that is often missing from training. Patients could benefit from more informed interpretation of ambiguous findings. Access to comparable prior cases may help reduce unnecessary work-up, missed escalation, delayed follow-up, or inappropriate reassurance. Hospital systems could transform imaging archives from passive storage into reusable institutional knowledge. Connecting images, reports, and downstream outcomes could support quality review, multidisciplinary discussion, trainee education, and automated teaching-file creation. Knowledge that already exists within the institution would become searchable and usable. Smaller and rural departments may benefit where subspecialty expertise is limited. Even a modest local archive can contain valuable precedent when it is visually indexed. An on-premise, read-only architecture with no required data egress could make the approach suitable for organizations where cloud-based tools are legally, technically, or operationally difficult.

GitHub Repo:
https://youtu.be/Bsaxpg4YjA8