

A novel AI-based pipeline for delivering radiologic education

The Atlas of Acute CT Brain Imaging — prototype

David Mark Rosewarne¹ · Mu'ath Ibrahim^{2,3} · Sahand Hooshmand^{2,3} · Niamh Moore⁴ · Noelle Clerkin · Seyedamir Tavakoli Taba³ · Ziba Gandomkar^{2,3} · Zhufeng Shi⁵ · Eric Myat · Jackie Chen · Tan Hasdarnakul⁶ · Anirudh Singh Bhati · Ella Grew · Diogo J. Vidal Silva⁷ · Patrick C. Brennan^{2,3}

¹ The Royal Wolverhampton NHS Trust, Wolverhampton, United Kingdom
² DetectedX, Sydney, Australia
³ The University of Sydney, Sydney, Australia
⁴ University College Cork, Cork, Ireland

⁵ The University of Queensland, Herston, Australia
⁶ Concord Repatriation General Hospital, Sydney, Australia
⁷ Norfolk and Norwich University Hospital NHS Foundation Trust, Norwich, United Kingdom

PURPOSE

Accessible, high-quality postgraduate radiology education is increasingly needed across lower-income countries, where growing numbers of trainee doctors are outpacing the time available to clinical educators. Two practical barriers stand in the way. The first is the **blank page problem**: the upfront cost of producing a well-structured educational series from scratch is prohibitively high for time-poor clinicians. The second is **case accessibility**: identifying suitable imaging examples within PACS archives containing tens of thousands of studies is a persistent challenge. This project applies AI to address both barriers, with a focus on **acute CT brain imaging**.

METHODS

Step 1.

Generating structured atlas text via nested LLM prompting.

Large language models such as GPT-5, with built-in web searching, allow garnering of text to user specification. We cannot expect a large educational document to be produced in a single pass with reliable results. We therefore imposed an atlas structure: **chapters** contain **pathologies**, and each pathology contains a fixed set of **subheadings**. **GPT-5** was used to generate the initial draft text at each leaf of this structure, and **Claude** was then used to revise, restructure and organise the output into the chapter-and-subchapter format used in the atlas, where each chapter covers one of the common pathologies in addition to the normal variants. Generation was delivered through **nested looping** in Python, making calls to the LLM APIs. The final prompt at each leaf of the loop took the form: “Give me lecture notes suitable for a radiology core trainee relating to the {subsection} of {pathology} in a chapter entitled {chapter}.” Draft material was then presented for **initial human editing** to verify text and remove hallucinations and other unhelpful material. The output subsequently underwent first-pass editing by trainee cohorts, followed by **consultant radiologist review**.

Step 2.

Linking generated text to real radiologic cases.

Suitable radiologic cases were sought to support the generated text. **Simple rule-based language processing**, implemented in Python, was used to find examples of the desired pathology from the **tens of thousands of CT brain examinations** residing on a PACS archive. These cases were embedded within the atlas for both **illustrative** and **assessment** purposes, so that each topic is anchored in real imaging and the assessment items the learner encounters are drawn from the same case base used to teach the topic.

RESULTS

We have generated an **Atlas of Acute CT Brain Imaging** through this AI-based educational pipeline. The authoring workflow runs as a guided **seven-stage process** on the **DetectedX** platform (Figure 1) — *Setup* → *Outline* → *Sections* → *Content* → *Enrich* → *Review* → *Publish* — with AI assistance available at each stage. Once published, atlas content is delivered to learners as interactive chapters combining annotated key images (Figure 2), inline knowledge checks and adaptive assessments (Figure 3), and the chapter-based reading experience shown in Figure 4.

We are currently in the **final phase of content development**, completing the chapters that cover the major acute CT brain pathologies and refining their assessments ahead of pilot deployment.

CONCLUSION

We have utilised AI as a tool for reducing the barrier to creating accessible and effective radiology education for young doctors in need. While the initial focus of this work is around **acute CT brain imaging**, once initial trials are conducted, this will be expanded to other key body regions, using the same **chapter-pathology-subheading structure** with the same **nested prompting and rule-based case retrieval**, so that it can be applied to any modality and organ system.

FUTURE WORK

Extension to other imaging domains. The pipeline is being adapted beyond brain CT to **chest** and **breast** imaging, using the same tooling for hierarchy generation, LLM prompting and PACS-linked case retrieval. **Trial in Vietnam, then expansion.** We aim to begin the trial of this material in Vietnam through **Hanoi Medical University**, and once complete, expand to other regions. **Learner-performance feedback.** Adaptive assessment data will be fed back into the authoring tool to surface weak topics and prioritise content for the next editorial cycle.

Figure 1. DetectedX authoring workflow — AI-assisted outline builder

The seven-stage workflow (Setup → Outline → Sections → Content → Enrich → Review → Publish) operationalises Methods Step 1. Here, the LLM-generated chapter hierarchy for *Haemorrhagic Brain Pathology* is open for editing.




Figure 2. Annotated sagittal CT — learner-facing atlas image

Atlas content is delivered to learners via the DetectedX platform. Each labelled structure is colour-coded by tissue class; sections of the atlas text are anchored to images such as this one.




Figure 3. DetectedX adaptive assessment — learner interface

Learners scroll the case stack and answer hierarchical questions on radiographic sign, anatomical localisation, and diagnosis. Cases retrieved via Step 2 populate the assessment bank; difficulty adapts to learner performance.




Figure 4. DetectedX learner view — published chapter on the platform

What learners see after publication: chapter prose, a key teaching point, an interactive viewer linking to scrollable annotated CT cases, and an inline knowledge check — all generated and curated through the pipeline shown in Figure 1.

