
CHAIR - CBM-Enabled Human-Al Collaboration for Image Retrieval

TECHNOLOGY NUMBER: 2025-048



OVERVIEW

Framework enables human correction of AI concepts for improved image retrieval accuracy

- Allows direct human intervention at concept level, enhancing interpretability and retrieval effectiveness
- Wildlife monitoring, medical imaging, educational image search, legal evidence analysis

BACKGROUND

Image retrieval systems are increasingly integral to fields like wildlife conservation and healthcare, where accurately locating specific images can enable conservation tracking or support medical diagnoses. Deep learning models have dramatically advanced retrieval capabilities, but struggle in complex real-world scenarios, often requiring human experts to review or rerun searches. Traditional "human-in-the-loop" approaches combine human and Al outputs post hoc, but the process can be cumbersome and inefficient due to limited interpretability within the Al models themselves. Humans are usually relegated to final decisions, unable to adjust or correct the model's conceptual understanding directly, which leads to lost time and missed opportunities for synergistic workflows. The need is clear for an approach that gives humans meaningful, efficient control within the image retrieval pipeline, leveraging both expertise and Al strengths.

Technology ID

2025-048

Category

Software

MOSS - Michigan Open Source Support

Inventor

Vaibhav Balloli Liz Bondi-Kelly Sara Beery

Further information

Ashwathi lyer ashwathi@umich.edu

View online page



CHAIR innovatively adapts the Concept Bottleneck Model (CBM) for image retrieval, creating an interface where humans can correct or refine intermediate concepts in the Al's processing pipeline. Unlike prior models that isolate the Al's internal reasoning, CHAIR's flexible intervention points allow users of varying expertise to adjust the underlying factors that drive image embeddings—improving interpretability, correctability, and ultimately, retrieval performance. This system achieved superior baseline metrics compared to similar models, even before human input. When expert intervention is incorporated, image retrieval accuracy improves further, demonstrating a new level of human-Al complementarity. Real-world applications span wildlife identification, medical case image search, legal image review, and educational content retrieval—any context where high-stakes or fine-grained image search benefits from expert insight and collaborative Al.

ADDITIONAL INFORMATION

PROJECT LINKS:

CHAIR Github

DEPARTMENT/LAB:

• Elizabeth Bondi-Kelly, Computer Science and Engineering

LICENSE:

MIT