

CBM-RAG: Demonstrating Enhanced Interpretability in Radiology Report Generation with Multi-Agent RAG and Concept Bottleneck Models

Hasan Md Tusfiqur Alam
hasan.alam@dfki.de

German Research Center for Artificial
Intelligence (DFKI)
Saarbrücken, Germany

Devansh Srivastav
devansh.srivastav@dfki.de

German Research Center for Artificial
Intelligence (DFKI)
Saarbrücken, Germany
Saarland University
Saarbrücken, Germany

Abdulrahman Mohamed Selim
abdulrahman.mohamed@dfki.de

German Research Center for Artificial
Intelligence (DFKI)
Saarbrücken, Germany

Md Abdul Kadir
abdul.kadir@dfki.de

German Research Center for Artificial
Intelligence (DFKI)
Saarbrücken, Germany
University of Oldenburg
Oldenburg, Germany

Md Moktadirul Hoque Shuvo
Dhaka Medical College Hospital
Dhaka, Bangladesh

Daniel Sonntag
daniel.sonntag@dfki.de
German Research Center for Artificial
Intelligence (DFKI)
Saarbrücken, Germany
University of Oldenburg
Oldenburg, Germany

Abstract

Advancements in generative Artificial Intelligence (AI) hold great promise for automating radiology workflows, yet challenges in interpretability and reliability hinder clinical adoption. This paper presents an automated radiology report generation framework that combines Concept Bottleneck Models (CBMs) with a Multi-Agent Retrieval-Augmented Generation (RAG) system to bridge AI performance with clinical explainability. CBMs map chest X-ray features to human-understandable clinical concepts, enabling transparent disease classification. Meanwhile, the RAG system integrates multi-agent collaboration and external knowledge to produce contextually rich, evidence-based reports. Our demonstration showcases the system's ability to deliver interpretable predictions, mitigate hallucinations, and generate high-quality, tailored reports with an interactive interface addressing accuracy, trust, and usability challenges. This framework provides a pathway to improving diagnostic consistency and empowering radiologists with actionable insights.

CCS Concepts

• **Applied computing** → **Health informatics**; • **Computing methodologies** → **Information extraction**; **Multi-agent systems**; • **Human-centered computing** → *Heat maps*; • **Information systems** → *Multimedia and multimodal retrieval*.

Keywords

Interpretable Radiology report generation, Disease classification, Medical imaging, Concept Bottleneck Models (CBM), Retrieval-Augmented Generation (RAG), Information Retrieval, VLMs, LLMs.

ACM Reference Format:

Hasan Md Tusfiqur Alam, Devansh Srivastav, Abdulrahman Mohamed Selim, Md Abdul Kadir, Md Moktadirul Hoque Shuvo, and Daniel Sonntag. 2025. CBM-RAG: Demonstrating Enhanced Interpretability in Radiology Report Generation with Multi-Agent RAG and Concept Bottleneck Models. In *Companion Proceedings of the 17th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS Companion '25)*, June 23–27, 2025, Trier, Germany. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3731406.3731970>

1 Introduction

Recent advancements in generative models have accelerated computer-aided interpretation for chest X-ray (CXR) images [2, 15, 18]. These end-to-end architectures not only predict specific findings but also generate comprehensive radiological reports by integrating a language module [9, 14]. A system that can classify diseases from CXR images and produce coherent reports can reduce radiologists' workload and improve diagnostic consistency. However, since large language models (LLMs) are prone to hallucinations [12], such generators face reliability issues. To address similar challenges in other domains, researchers have introduced Retrieval-Augmented Generation (RAG) [8], which leverages external resources to produce more accurate and reliable conclusions. However, the black-box nature of LLMs remains a significant limitation [10], as they fail to provide explanations or interpretable relationships between inputs and outputs, leading to a system that may be perceived as unreliable and untrustworthy. Trust in these systems requires transparency [4], interpretability [6], and integration of additional data such as patient history and recent research.

To address these challenges, we propose a conversational tool integrating Concept Bottleneck Models (CBMs) [7] with a multi-agent RAG framework to enhance accuracy, interpretability, and reliability in CXR report generation. CBMs map visual features to human-understandable clinical concepts and use saliency techniques to highlight relevant image regions, while RAG dynamically incorporates external knowledge, including patient history, prior studies, and current research, to produce evidence-based reports. In this paper, we demonstrate an end-to-end implementation that combines interpretable disease classification with robust report generation, mitigating issues of hallucination and opacity, and thereby enhancing AI-driven CXR interpretation for clinical practice to empower radiologists with actionable insights to improve their diagnostic consistency and trust in our system. Our code¹ and demo² are publicly available online.

2 Methodology

Our approach starts with a concept bottleneck mechanism [7] to identify and quantify medically relevant concepts in a CXR image. Building on prior works [1, 11, 16], we use LLMs to automatically acquire a set of concepts for classification, rather than relying on manual identification. As shown in Fig. 1, we obtain image embedding and text embeddings for the uploaded image and concept set, respectively, from ChexAgent [2], a VLM fine-tuned for CXR interpretation, and the Mistral embed model[5]. We calculate cosine similarity between image embeddings and each text embedding in the concept set to form a similarity matrix. To focus on the most significant features, max pooling is applied to the similarity matrix to form a concept vector. This concept vector is normalized to a scale between 0 and 1 for interpretability and fed into a fully connected layer. A classification model uses this vector as input to predict the disease class. We used the COVID-QU dataset [3], comprising 33,920 CXR images with three classes: Pneumonia, COVID-19, and Normal. Finally, the cross-product of the model's weight matrix and the concept vector provides contribution scores, quantifying the influence of each concept on the classification decision. Saliency regions for each concept in the image are derived from the similarity matrix. These heatmaps serve as direct visual indications of how the system localizes concepts such as “pulmonary consolidation” or “nodule” within the CXR image, thereby offering a clear route to interpretability.

In addition, we used a multi-agent RAG with five specialized agents for report generation. The Pneumonia, COVID-19, and Normal Agents are implemented as Reasoning and Acting (ReAct) agents [17]. Additionally, the Radiologist Agent interprets clinical concepts using the ReAct agents and queries a pre-configured database from the National Institutes of Health (NIH), while the Report Writer Agent synthesizes the final report. The system also accepts user-provided files (e.g., PDFs, PPTs, text, MP3, MP4), with media transcribed via OpenAI's Whisper model [13] and embedded and indexed for retrieval. This integration enriches reports with updated clinical guidelines, patient histories, and multimedia sources. The framework is implemented using CrewAI and LlamaIndex for efficient retrieval and high-quality report generation.




Figure 1: Workflow of the CBM-RAG Framework for Radiology Report Generation. The upper section processes chest X-rays via a VLM to generate clinical concepts, heatmaps, and contribution scores. The lower section uses multi-agent RAG. A Radiologist Agent synthesizes findings, a Report Writer Agent creates detailed reports, and a Chat Agent enables real-time interaction.

3 User Interface

The user interface (UI) for the CXR analysis system comprises three components: concept generation, report generation, and a conversational chat interface. Upon uploading a CXR image, the concept generation module identifies relevant clinical concepts, computes contribution scores, and predicts disease classes using the CBM. Identified concepts are displayed in an editable list sorted by the absolute values of their contribution scores, each with a toggle for visualizing associated saliency heatmaps. Users can adjust scores to refine model predictions, thereby linking outputs to clinically meaningful features. After finalizing concept scores, users can generate a comprehensive radiology report. The report generation module integrates clinical documents from trusted sources (e.g., NIH) and accepts additional inputs (text, audio, video, images). The generated report details findings, diagnosis, and guidelines, and an optional chain-of-thought dropdown reveals the multi-agent RAG's sequential reasoning. A conversational chat interface further enables real-time, context-aware queries regarding the CXR image, report details, or clinical conditions.

4 Conclusion and Future Work

In this paper, we presented a tool that bridges AI performance with clinical explainability by linking visual features to human-understandable clinical concepts and integrating external knowledge for context-rich, evidence-based radiology reports. Our framework produces transparent disease classifications and tailored reports while mitigating hallucination and opacity issues. Its interactive UI—with explainable outputs and conversational capabilities—facilitates dynamic clinician engagement, enhancing trust in AI-assisted decision-making. Although technically promising, formal usability studies in real clinical settings are yet to be conducted. Future work will include comprehensive user evaluations, extension to other imaging modalities, and exploration of broader healthcare applications.

¹Code: <https://github.com/tifat58/enhanced-interpretable-report-generation-demo.git>

²Online Demo: <https://cxr-cbm-rag-dfki-impl-demo.streamlit.app/>

Acknowledgments

This work was funded by the German Federal Ministry of Education and Research (BMBF) under grant number 01IW23002 (No-IDLE) and by the Endowed Chair of Applied AI at the University of Oldenburg.

References

- [1] Hasan Md Tusfiqur Alam, Devansh Srivastav, Md Abdul Kadir, and Daniel Sonntag. 2025. Towards Interpretable Radiology Report Generation via Concept Bottlenecks Using a Multi-agent RAG. In *Advances in Information Retrieval - 47th European Conference on Information Retrieval, ECIR 2025, Lucca, Italy, April 6-10, 2025, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 15574)*, Claudia Hauff, Craig Macdonald, Dietmar Jannach, Gabriella Kazai, Franco Maria Nardini, Fabio Pinelli, Fabrizio Silvestri, and Nicola Tonello (Eds.). Springer, 201–209. doi:10.1007/978-3-031-88714-7_18
- [2] Zhihong Chen, Maya Varma, Jean-Benoit Delbrouck, Magdalini Paschali, Louis Blankemeier, Dave Van Veen, Jeya Maria Jose Valanarasu, Alaa Youssef, Joseph Paul Cohen, Eduardo Pontes Reis, Emily Tsai, Andrew Johnston, Cameron Olsen, Tanishq Mathew Abraham, Sergios Gatidis, Akshay S Chaudhari, and Curtis Langlotz. 2024. CheXagent: Towards a Foundation Model for Chest X-Ray Interpretation. In *AAAI 2024 Spring Symposium on Clinical Foundation Models*. <https://openreview.net/forum?id=P3LOmrZWGR>
- [3] Muhammad EH Chowdhury, Tawsifur Rahman, Amith Khandakar, Rashid Mazhar, Muhammad Abdul Kadir, Zaid Bin Mahbub, Khandakar Reajul Islam, Muhammad Salman Khan, Atif Iqbal, Nasser Al Emadi, et al. 2020. Can AI help in screening viral and COVID-19 pneumonia? *Ieee Access* 8 (2020), 132665–132676.
- [4] Christopher Ifeanyi Eke and Liyana Shuib. 2024. The role of explainability and transparency in fostering trust in AI healthcare systems: a systematic literature review, open issues and potential solutions. *Neural Computing and Applications* (2024), 1–36.
- [5] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).
- [6] Ujjwal Singh Kathait, Anamika Rana, Rahul Chauhan, and Ruchira Rawat. 2024. A Comprehensive Review of Interpretability in AI and Its Implications for Trust in Critical Applications. In *2024 4th International Conference on Sustainable Expert Systems (ICES)*. IEEE, 1683–1693.
- [7] Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept bottleneck models. In *International conference on machine learning*. PMLR, 5338–5348.
- [8] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [9] Zhengliang Liu, Aoxiao Zhong, Yiwei Li, Longtao Yang, Chao Ju, Zihao Wu, Chong Ma, Peng Shu, Cheng Chen, Sekeun Kim, et al. 2023. Tailoring large language models to radiology: A preliminary approach to llm adaptation for a highly specialized domain. In *International Workshop on Machine Learning in Medical Imaging*. Springer, 464–473.
- [10] Haoyan Luo and Lucia Specia. 2024. From understanding to utilization: A survey on explainability for large language models. *arXiv preprint arXiv:2401.12874* (2024).
- [11] Tuomas Oikarinen, Subhro Das, Lam M Nguyen, and Tsui-Wei Weng. 2023. Label-free concept bottleneck models. *arXiv preprint arXiv:2304.06129* (2023).
- [12] Gabrijela Perković, Antun Drobniak, and Ivica Botički. 2024. Hallucinations in llms: Understanding and addressing challenges. In *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*. IEEE, 2084–2088.
- [13] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*. PMLR, 28492–28518.
- [14] Zhanyu Wang, Lingqiao Liu, Lei Wang, and Luping Zhou. 2023. R2gengpt: Radiology report generation with frozen llms. *Meta-Radiology* 1, 3 (2023), 100033.
- [15] Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. 2022. MedCLIP: Contrastive Learning from Unpaired Medical Images and Text. *arXiv:2210.10163 [cs.CV]* <https://arxiv.org/abs/2210.10163>
- [16] An Yan, Yu Wang, Yiwu Zhong, Zexue He, Petros Karypis, Zihan Wang, Chengyu Dong, Amilcare Gentili, Chun-Nan Hsu, Jingbo Shang, et al. 2023. Robust and interpretable medical image classifiers via concept bottleneck models. *arXiv preprint arXiv:2310.03182* (2023).
- [17] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629* (2022).
- [18] Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, Andrea

Tupini, Yu Wang, Matt Mazzola, Swadheen Shukla, Lars Liden, Jianfeng Gao, Matthew P. Lungren, Tristan Naumann, Sheng Wang, and Hoifung Poon. 2024. BiomedCLIP: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. *arXiv:2303.00915 [cs.CV]* <https://arxiv.org/abs/2303.00915>