MedHyperGraph: EHR-Integrated Multimodal Hyperedges for Clinical VQA

Mina Heinein[1,⋆], Mina Youssef[1,⋆], Kirellos Nashed[1,⋆], Tamer Basha[1], Hasan Md Tusfiqur Alam[2], Abdulrahman Mohamed Selim[2], Omair Shahzad Bhatti[(✉)2], and Daniel Sonntag[2,3]

[1] Faculty of Engineering, Cairo University, Egypt
[2] German Research Center for Artificial Intelligence (DFKI), Germany
[3] University of Oldenburg, Germany
(✉)omair_shahzad.bhatti@dfki.de

Abstract. Many clinical questions about chest X-rays cannot be answered from images alone, as they depend on longitudinal context, including prior studies, earlier reports, diagnoses, and the timing of clinical events. Most medical vision-language models still process isolated image-question pairs. We present MedHyperGraph a per-patient multimodal hypergraph for electronic health record (EHR)-grounded medical visual question answering (VQA). It integrates structured EHR events, prior-report entities, imaging studies, ontology concepts, and image regions into a single temporally indexed structure; evidence is then retrieved under explicit temporal constraints. On EHRXQA boolean questions ($n = 1{,}900$), the retrieval baselines remain close to a majority-class predictor on the MedGemma backbones; MedHyperGraph improves on all three backbones, reaching 73.95% accuracy (95% CI $[71.9, 75.9]$), without task-specific training. On open-ended questions, it raises weighted F1 by 10.0 points at 27B, driven by multi-study questions ($5.2 \rightarrow 43.4$). Ablations attribute the gain primarily to the structured EHR (-9.1) indicating that EHR-grounded VQA benefits from retrieval that preserves the temporal and n-ary structure of clinical evidence. Code on GitHub.

Keywords: Medical VQA · Electronic health records · Multimodal retrieval · Hypergraphs · Retrieval-augmented generation · Knowledge Graph

1 Introduction and Related Work

A chest X-ray is rarely interpreted in isolation, because its meaning depends on the patient's history, including prior admissions, studies, reports, diagnoses, labs, and the time between events [3,24,2]. Whether a finding is new, unchanged, expected after a procedure, or consistent with a known diagnosis is often not visible in a single image. Medical visual question answering (VQA) should reflect this, because answers may depend on when the image was taken, which clinical event it follows, or whether a related condition already appears in the record.

⋆ These authors contributed equally to this work.

Recent medical vision-language models (VLMs) such as MedGemma [19], LLaVA-Med [13], and Lingshu [20] interpret medical images well, but typically answer isolated image-question pairs with little access to the surrounding record. This is a poor fit for EHRXQA [3], a benchmark over EHR tables and chest X-rays whose questions can combine structured SQL queries with visual reasoning over MIMIC-CXR images [11]; many of its questions cannot be answered from the image alone and need evidence from across the patient's longitudinal record.

Retrieval-augmented generation (RAG) can provide evidence to a VLM [9], but flat text retrieval does not capture the temporal and multimodal relations that clinical reasoning relies on. Graph-based RAG represents knowledge as entities and relations [8], but a single clinical fact usually combines several elements at once (e.g., a finding, an anatomical region, a study time, a report statement, and an EHR event); pairwise edges cannot represent this grouping unless the fact is represented explicitly. HyperGraphRAG [16] preserves such n-ary facts through hyperedges, as do recent claim-centric knowledge-graph approaches for multimodal personal data [5]. In medicine, multimodal knowledge graphs such as MEDMKG [22], VLM-KG [1], and R2GenKG [21] link images to clinical concepts; however, they do not target patient-specific longitudinal retrieval for EHR-grounded VQA, where evidence must be jointly selected from EHR events, prior reports, timestamps, and image regions. Recently, Yang et al. [24] introduced HyperWalker, which builds a dynamic clinical hypergraph and trains a reinforcement-learning agent to traverse diagnostic paths, achieving 70.43% boolean accuracy on a filtered EHRXQA test set. Although HyperWalker shows the value of hypergraph structures on this benchmark, it learns a global navigation policy, does not isolate each patient's own evidence, and does not constrain what the model may read at answer time. This motivates a simpler, patient-scoped formulation that needs no training and keeps its evidence explicit.

Therefore, we present MedHyperGraph, a per-patient multimodal hypergraph for EHR-grounded medical VQA. For each patient, MedHyperGraph integrates structured EHR events, prior-report entities, imaging studies, ontology concepts, and image regions into a single structure; cross-modal links are computed from BioViL-T [6] embeddings; the hypergraph design is informed by RadGraph [10], Chest ImaGenome [23], and CheXRelNet [12]. Given a question, MedHyperGraph retrieves a small evidence subgraph under explicit temporal constraints and passes it, together with the current image, to a VLM. The target study's own report is withheld from retrieval, and future studies are excluded, so a visual answer cannot be copied from the current report. Unlike HyperWalker, MedHyperGraph is patient-scoped rather than global, does not use task-specific training, and selects evidence by retrieval rather than a learned policy.

Our Contributions. We introduce MedHyperGraph, a per-patient multimodal hypergraph that unifies EHR events, prior-report entities, imaging studies, ontology concepts, and image regions for EHR-grounded medical VQA. Additionally, we propose temporally aware retrieval with leakage control, ontology-grounded normalization with SNOMED CT [7], and an evaluation on the filtered EHRXQA

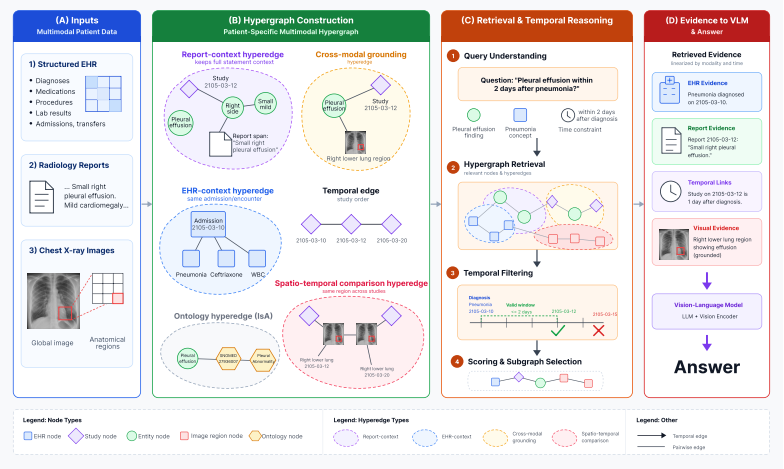


Fig. 1. Overview of MedHyperGraph. Multimodal patient data from EHRs, radiology reports, and chest X-rays form a patient-specific hypergraph, retrieved with temporal constraints, and linearized as evidence for VLM-based answer generation.

test set against vanilla inference, text-RAG, and GraphRAG baselines, including temporal and clinical-attribute breakdowns. Our code is available on GitHub.

2 Methodology

MedHyperGraph represents each patient record as a typed multimodal hypergraph rather than as isolated image-question samples. For patient p, we define $\mathcal{G}_p = (\mathcal{V}_p, \mathcal{E}_p)$, where $\mathcal{V}_p$ is the set of clinical hypernodes and $\mathcal{E}_p$ is the set of typed hyperedges. Each hypernode represents a clinically meaningful unit, such as an EHR event, radiology study, report entity, image region, or ontology concept. Each hyperedge is defined as $e = (S_e, \rho_e, \alpha_e)$ with $S_e \subseteq \mathcal{V}_p$ and $|S_e| \geq 2$, where ρ_e denotes the hyperedge type and α_e stores attributes such as timestamps, source text, grounding scores, or ontology identifiers. Unlike a standard graph edge, a hyperedge can connect more than two clinical elements, allowing MedHyperGraph to preserve higher-order clinical contexts such as report statements, admission events, cross-modal image-text links, and longitudinal comparisons.

2.1 Problem Formulation

For each patient p, let $\mathcal{D}_p = (\mathcal{T}_p, \mathcal{R}_p, \mathcal{I}_p)$ denote the longitudinal record, consisting of structured EHR tables $\mathcal{T}_p$, radiology reports $\mathcal{R}_p$, and chest X-ray studies $\mathcal{I}_p$. Each study $s \in \mathcal{I}_p$ has a timestamp $\tau(s)$, one or more images, and an associated report. Given a question $x = (p, q)$, where q is a natural-language clinical question about patient p, the task is to predict $\hat{y} \in \{\text{yes}, \text{no}\}$.

At inference time, MedHyperGraph retrieves an evidence subgraph $\mathcal{S}_q = (\mathcal{V}_q, \mathcal{E}_q)$, with $\mathcal{V}_q \subseteq \mathcal{V}_p$ and $\mathcal{E}_q \subseteq \mathcal{E}_p$, and conditions a VLM on textual and visual evidence (Figure 1):

$$\hat{y} = F_\theta\big(q,\, \mathcal{L}(\mathcal{S}_q),\, \mathcal{M}(\mathcal{S}_q)\big).$$

Here, $\mathcal{L}(\mathcal{S}_q)$ is the textual linearization of the retrieved evidence, and $\mathcal{M}(\mathcal{S}_q)$ contains retrieved visual evidence such as images or anatomical regions.

2.2 MedHyperGraph Construction

For each patient, MedHyperGraph builds a typed hypernode set $\mathcal{V}_p = \mathcal{V}_p^{\text{ehr}} \cup \mathcal{V}_p^{\text{study}} \cup \mathcal{V}_p^{\text{ent}} \cup \mathcal{V}_p^{\text{img}} \cup \mathcal{V}_p^{\text{ont}}$. **EHR hypernodes** $\mathcal{V}_p^{\text{ehr}}$ are derived from structured tables, including admissions, diagnoses, procedures, prescriptions, laboratory events, and transfers. **Study hypernodes** $\mathcal{V}_p^{\text{study}}$ represent radiology studies and serve as temporal anchors linking images, reports, and nearby EHR evidence. **Report-entity hypernodes** $\mathcal{V}_p^{\text{ent}}$ are extracted from reports and include findings, anatomical regions, severity modifiers, devices, and other clinical concepts; when possible, mentions are normalized to SNOMED CT identifiers [7]. **Image hypernodes** $\mathcal{V}_p^{\text{img}}$ represent chest X-rays at both global-image and anatomical-region levels. **Ontology hypernodes** $\mathcal{V}_p^{\text{ont}}$ encode normalized concepts and their parent concepts to support synonym-aware expansion.

MedHyperGraph defines hyperedges as a union of typed clinical relations:

$$\mathcal{E}_p = \mathcal{E}_p^{\text{str}} \cup \mathcal{E}_p^{\text{rep}} \cup \mathcal{E}_p^{\text{ehr}} \cup \mathcal{E}_p^{\text{cmp}} \cup \mathcal{E}_p^{\text{ont}}.$$

Structural hyperedges $\mathcal{E}_p^{\text{str}}$ link naturally paired elements, such as studies and their reports or images, images and anatomical regions, or consecutive studies. **Report-context hyperedges** $\mathcal{E}_p^{\text{rep}}$ preserve statement-level meaning by linking a study node to all report entities extracted from the same text segment. For a segment c associated with study s, we define $e_c^{\text{rep}} = (\{s\} \cup \text{Ent}(c), \text{ReportContext}, \alpha_c)$, where α_c stores the original text span. **EHR-context hyperedges** $\mathcal{E}_p^{\text{ehr}}$ link an admission node to structured events in the same clinical context, e.g., diagnoses, prescriptions, procedures, and lab results. **Spatio-temporal comparison hyperedges** $\mathcal{E}_p^{\text{cmp}}$ connect corresponding anatomical regions across consecutive studies for longitudinal reasoning. **Ontology hyperedges** $\mathcal{E}_p^{\text{ont}}$ connect entity mentions to normalized concepts and their parents.

Representational motivation. Hyperedges preserve statement-level clinical context that a naive pairwise graph can lose: a joint fact over a finding, region, timestamp, and EHR event can collapse to the same clique as independent pairwise links. Storing these elements as a single hyperedge preserves shared membership and lets MedHyperGraph recover the full local context by incidence expansion from a partial retrieval hit, before ranking and context-budget pruning.

2.3 Temporally Aware Multimodal Retrieval

Given a question q, the retrieval module selects patient-specific evidence from $\mathcal{G}_p$ in four steps (Figure 1). **First**, the question is analyzed to identify clinical

concepts, referenced studies, anatomical regions, and temporal constraints. The query is encoded as $z_q = f_{\text{text}}(q)$; when a study or image is explicitly referenced, the corresponding image representation $z_I = f_{\text{img}}(I)$ is also used. Temporal expressions define a window $W(q)$ relative to an admission, diagnosis, or study.

Second, candidate evidence is retrieved through three complementary channels: (i) entity-level retrieval over EHR, report-entity, and ontology hypernodes; (ii) hyperedge-level retrieval over complete clinical contexts; and (iii) visual retrieval over image or anatomical-region hypernodes when available.

Third, candidates are expanded through hypergraph incidence links. If a retrieved hypernode belongs to a hyperedge, the full hyperedge and its incident hypernodes are added. If a hyperedge is retrieved directly, all of its connected hypernodes are included. Candidates inconsistent with the temporal window $W(q)$ are removed or down-weighted. **Finally**, each candidate u is scored using textual and visual relevance, temporal compatibility, and hypergraph proximity:

$$S(u, q) = \lambda_{\text{text}} s_{\text{text}}(u, q) + \lambda_{\text{vis}} s_{\text{vis}}(u, q) + \lambda_{\tau} s_{\tau}(u, q) + \lambda_h s_h(u, q). \quad (1)$$

Equation (1) describes the signal space; the reported system enables or disables each term rather than tuning continuous weights, which we leave to future work. The top-scoring hypernodes and hyperedges under a fixed context budget form the final evidence subgraph $\mathcal{S}_q$, which is linearized by temporal order into three blocks: *Entities*, *Relationships* (each hyperedge with its complete member set), and *Sources* (reports and EHR text). The current study's images are attached to the same call, and the VLM generates the final yes/no or short answer.

3 Experimental Setup

Dataset and metric. We evaluate on EHRXQA [3], built over MIMIC-IV and MIMIC-CXR [11], using the HyperWalker-filtered test set [24]: 1,900 boolean and 600 open-ended questions. The attribute analysis uses the 719 boolean questions that require reading the image. The boolean answers are imbalanced (56% "no", 44% "yes"), so a constant "no" already scores 56.0% accuracy at 50.0% balanced accuracy. Therefore, we report *balanced accuracy* as the primary metric. Dates in MIMIC are shifted into the future during de-identification, so years such as "2105" are expected rather than errors.

Backbones and baselines. We use three backbones: MedGemma 1.5 (4B), MedGemma 1 (27B), and Qwen 3.6-35B-A3B-AWQ [4]. Under each, we run *Vanilla* (no retrieval), *text-RAG* (dense retrieval over prior-report and EHR text), *GraphRAG* [8] (a naive pairwise entity graph over the same text), and *MedHyperGraph*. All retrieval methods share the same text sources and context budget. Lingshu [20] and HyperWalker [24] use their own backbones and training and serve only as external reference points.

Leakage control. Many EHRXQA visual answers also appear in the study's own radiology report, so we control what the model may read at answer time. Retrieval is restricted to data dated before the query time, so future studies are never used. The target study's own report is withheld from the retrievable pool,

so a visual answer cannot be copied from the current report; only prior reports are retrievable. Image regions use Chest ImaGenome [23] bounding-box geometry only, not its finding labels, so report-derived labels do not enter retrieval. Retrieval weights are fixed in advance and are not tuned on the test set.

Open-ended scoring. n-gram metrics such as BLEU [18] and ROUGE [14] penalize clinically equivalent paraphrases, while BERTScore [26] does not check factual correctness [17]. Therefore, we use two LLM judges [27,15]: Gemma-4-31B and Qwen 3.6-35B-A3B-AWQ [4]; both run locally with deterministic decoding, a fixed rubric prompt, and structured JSON output; we report their average. Following ReXVal [25] and GREEN [17], each judge identifies matched, missed, and hallucinated findings before scoring recall, precision, correctness, and semantic similarity on a 0–100 scale. We report F1 as the harmonic mean of mean recall and mean precision. We evaluate open-ended questions on all three backbones and interpret within-backbone differences, since a judge sharing a model family with a generator may favor its own outputs. Additional implementation details are available on our GitHub repository.

4 Results and Discussion

Boolean accuracy (Table 1). On both MedGemma backbones, the retrieval baselines stay close to the majority-class predictor: 58–59% accuracy, only 2–3 points above the 56% of a constant "no" predictor, and 54.8–58.1% balanced accuracy, only 5–8 points above chance. Flat text retrieval and a binary entity graph add little over no retrieval. This ceiling is not universal; with the stronger Qwen backbone, the baselines recover (66.66% balanced accuracy for text-RAG), indicating that the effect depends on backbone capability rather than on a weak baseline implementation. MedHyperGraph performs best on every backbone and is the only method that improves regardless of backbone, reaching 63.61%, 69.13%, and 73.64% balanced accuracies. It also nearly closes the accuracy–balanced-accuracy gap on Qwen (73.95 vs. 73.64), so it discriminates both answers rather than favoring "no".

Open-ended VQA (Table 1). MedHyperGraph achieves the best overall F1 across all three backbones and is the only method that improves over the vanilla model at every one: MedGemma +3.8 at 4B and +10.0 at 27B, and +11.8 on Qwen. GraphRAG lowers F1 at both MedGemma backbones(−4.6 at 4B, −7.7 at 27B) and adds little on Qwen (+2.0), while text-RAG helps at 4B (+3.5) and on Qwen (+9.7) but hurts at 27B (−6.5). The gain comes from multi-study (N-image) questions, where vanilla models nearly fail (F1 5.2 at both MedGemma sizes, 17.3 on Qwen) and MedHyperGraph reaches 19.0 and 43.4 and 54.1. The advantage is not uniform: at 4B text-RAG is stronger on single- and two-image questions, and at 27B, the vanilla model keeps a small edge on two-image questions (32.1 vs. 29.7). MedHyperGraph wins overall through its large multi-study margin, where longitudinal evidence is decisive.

Component ablation (Table 2). Removing evidence sources from the system shows where the accuracy comes from. Structured EHR evidence contributes

Table 1. Results on the filtered EHRXQA test set. Boolean questions (56% "no") report accuracy and balanced accuracy, the primary metric, with 95% Wilson intervals. Open-ended questions report weighted F1 by image count, averaged over two LLM judges (Gemma-4-31B and Qwen 3.6). Best method per backbone and metric in **bold**. HyperWalker (RL-trained, own backbone) reports 70.43% boolean accuracy.

	Boolean (n = 1,900)				**Open-ended F1** (n = 600)			
Method	Acc.	95% CI	Bal. Acc.	95% CI	1-img	2-img	N-img	Overall
MedGemma 1.5 (4B)								
Vanilla	59.26	[57.04, 61.45]	54.99	[53.30, 56.71]	33.4	26.6	5.2	22.9
text-RAG	58.47	[56.23, 60.67]	54.84	[52.97, 56.73]	**38.5**	**28.1**	11.7	26.4
GraphRAG [8]	58.74	[56.51, 60.93]	54.88	[53.07, 56.72]	28.6	15.8	10.0	18.3
MedHyperGraph	**64.63**	[62.45, 66.75]	**63.61**	[61.42, 65.74]	33.5	25.9	**19.0**	**26.7**
MedGemma 1 (27B)								
Vanilla	58.11	[55.87, 60.31]	56.54	[54.36, 58.71]	43.4	**32.1**	5.2	28.6
text-RAG	58.94	[56.72, 61.14]	58.13	[55.90, 60.33]	26.5	16.1	25.4	22.1
GraphRAG [8]	58.84	[56.61, 61.03]	54.92	[53.13, 56.74]	35.5	19.0	4.3	20.9
MedHyperGraph	**69.40**	[67.29, 71.43]	**69.13**	[66.98, 71.17]	**45.3**	29.7	**43.4**	**38.6**
Qwen 3.6								
Vanilla	60.00	[57.78, 62.18]	57.67	[55.58, 59.75]	43.4	33.1	17.3	32.4
text-RAG	67.89	[65.76, 69.96]	66.66	[64.51, 68.73]	**49.1**	**41.7**	34.8	42.1
GraphRAG [8]	64.84	[62.67, 66.96]	62.40	[60.33, 64.43]	48.2	34.1	19.7	34.4
MedHyperGraph	**73.95**	[71.93, 75.87]	**73.64**	[71.61, 75.57]	48.9	31.9	**54.1**	**44.2**

Table 2. Component ablation (accuracy with 95% Wilson intervals, n = 1,900).

	MedGemma 27B			**Qwen 3.6**		
Configuration	Acc.	95% CI	Δ	Acc.	95% CI	Δ
MedHyperGraph (full)	**69.40**	[67.29, 71.43]	–	**73.95**	[71.93, 75.87]	–
w/o EHR evidence	60.35	[58.13, 62.53]	−9.05	64.21	[62.03, 66.33]	−9.74
w/o prior reports	62.57	[60.37, 64.72]	−6.83	71.21	[69.13, 73.20]	−2.74
w/o temporal filter	68.50	[66.38, 70.55]	−0.90	73.79	[71.77, 75.72]	−0.16
w/ grounding	67.70	[65.56, 69.77]	−1.70	70.41	[68.32, 72.42]	−3.54

most (−9.05 and −9.74 when removed) and prior reports second (−6.83 and −2.74), which supports the central claim that longitudinal, multi-source evidence drives the gain. The prior-report effect is smaller on Qwen, suggesting that a stronger backbone recovers more from the EHR context alone. Temporal filtering contributes −0.90 and −0.16; this small effect is expected, since the filter is inactive for questions without a time constraint, and its benefit concentrates on the time-windowed subset rather than the full test set. Unlike the removal ablations, we test visual grounding as an optional addition. The base system already passes the image to the VLM, while grounding adds inferred entity-region links. Adding them lowers accuracy on both backbones, suggesting that similarity-based alignment is too noisy and distracts retrieval. Although it may support visual inspection, improving alignment remains future work.

Temporal reasoning. Grouping the 27B questions by their temporal form (each question belongs to one category) shows where the temporal filter helps. On the two time-windowed categories the baselines stay at or below the vanilla

model: global-window, base 57.0 vs. 53.7 (text-RAG) and 54.5 (GraphRAG); relative-window ("within"), 52.3 vs. 57.3 and 48.6; in contrast, MedHyperGraph reaches 71.5 and 66.7 balanced accuracy. The same ordering holds at 4B (53.2 → 67.3 and 52.9 → 60.6). Non-temporal questions also improve (55.8 → 77.0 at 27B), but that gain comes from structural and multimodal retrieval rather than the time window, which is consistent with the small overall ablation effect above. Change-comparison is the hardest category, where MedHyperGraph (60.3) trails text-RAG (63.9); these questions ask about progression across studies and remain the clearest target for future work.

Analysis by clinical attribute. Improvement varies across finding types. On the 719 visual questions, the gain over the vanilla 27B model is largest for findings that are named explicitly in prior reports and correspond to discrete structures, such as support devices and lines (+24.4 balanced accuracy) and focal opacities (+21.1). It is smaller for diffuse parenchymal findings, where the vanilla model already performs comparatively well, for example, chronic lung disease (+3.5); this follows from the retrieval mechanism, where report entities that carry a stable name and a consistent anatomical location are normalized to ontology concepts and grounded to image regions more reliably than findings described in varying terms. Two attribute groups depart from this pattern; for pleural and air-related findings, text-RAG reaches 73.0 against 68.1 for MedHyperGraph, which is plausible because such findings are typically stated in plain language in the report and are retrievable without additional structure. The *technical* group contains only 12 questions and does not support interpretation. Group sizes range from 12 to 122, and the comparisons are not corrected for multiple testing, so these differences indicate tendencies rather than established effects.

Results Summary. Across boolean and open-ended questions, and temporal and attribute breakdowns, flat and binary-graph retrieval are inconsistent and in several settings degrade the vanilla model, while MedHyperGraph improves on all three backbones. MedHyperGraph's advantage lies in questions whose answers depend on longitudinal evidence, namely multi-study and time-windowed questions; the ablation attributes the gains primarily to the structured EHR.

5 Conclusion and Future Work

We presented MedHyperGraph, a patient-specific multimodal hypergraph that unifies EHR events, prior reports, imaging studies, ontology concepts, and image regions into a temporally indexed retrieval structure for EHR-grounded medical VQA. On EHRXQA, MedHyperGraph outperformed vanilla inference, text-RAG, and GraphRAG on all three backbones without task-specific training, and improved open-ended weighted F1, driven mainly by multi-study questions. Ablations attribute most of the gain to structured EHR. Gains were clearest for longitudinal queries, while cross-study change-comparison remained the hardest category, a target for future work such as agentic question decomposition. We also built a simple UI, shown in the supplementary video, for future user studies.

Acknowledgments. This work was partially funded by the Federal Ministry of Research, Technology and Space (BMFTR) under grant number 16IW23002 (No-IDLE), grant number 16IW24006 (NoIDLEChatGPT); the Lower Saxony Ministry of Science and Culture (MWK) in the zukunft.niedersachsen program; and the Endowed Chair of AAI at the University of Oldenburg.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdullah, A., et al.: VLM-KG: Multimodal Radiology Knowledge Graph Generation (May 2025). https://doi.org/10.48550/arXiv.2505.17042
2. Alam, H.M.T., et al.: Towards interpretable radiology report generation via concept bottlenecks using a multi-agentic RAG. In: Advances in Information Retrieval - European Conference on Information Retrieval, ECIR. Lecture Notes in Computer Science, Springer (2025). https://doi.org/10.1007/978-3-031-88714-7_18
3. Bae, S., et al.: EHRXQA: a multi-modal question answering dataset for electronic health records with chest X-ray images. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. pp. 3867–3880. NIPS '23, Curran Associates Inc., Red Hook, NY, USA (Dec 2023), https://openreview.net/forum?id=Pk2x7FPuZ4
4. Bai, S., et al.: Qwen3-VL Technical Report (Nov 2025). https://doi.org/10.48550/arXiv.2511.21631
5. Bhatti, O.S., et al.: Photograph: Claim-centric knowledge graphs for personal photo management. In: Proceedings of the GenAIK-NORA Joint Workshop, IJCAI-ECAI (2026)
6. Boecking, B., et al.: Making the Most of Text Semantics to Improve Biomedical Vision–Language Processing. In: Computer Vision – ECCV 2022. pp. 1–21. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-20059-5_1
7. Donnelly, K.: SNOMED-CT: The Advanced Terminology and Coding System for eHealth. In: Medical and Care Compunetics 3, pp. 279–290. IOS Press (2006), https://ebooks.iospress.nl/publication/9130
8. Edge, D., et al.: From Local to Global: A Graph RAG Approach to Query-Focused Summarization (Feb 2025). https://doi.org/10.48550/arXiv.2404.16130
9. Gao, Y., et al.: Retrieval-Augmented Generation for Large Language Models: A Survey (Mar 2024). https://doi.org/10.48550/arXiv.2312.10997
10. Jain, S., et al.: RadGraph: Extracting Clinical Entities and Relations from Radiology Reports. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Jun 2021), https://openreview.net/forum?id=pMWtc5NKd7V
11. Johnson, A.E.W., et al.: MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. Scientific Data **6**(1), 317 (Dec 2019). https://doi.org/10.1038/s41597-019-0322-0
12. Karwande, G., et al.: CheXRelNet: An Anatomy-Aware Model for Tracking Longitudinal Relationships Between Chest X-Rays. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 581–591. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-16431-6_55

13. Li, C., et al.: LLaVA-med: training a large language-and-vision assistant for biomedicine in one day. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. pp. 28541–28564. NIPS '23, Curran Associates Inc., Red Hook, NY, USA (Dec 2023), https://dl.acm.org/doi/10.5555/3666122.3667362
14. Lin, C.Y.: ROUGE: A Package for Automatic Evaluation of Summaries. In: Text Summarization Branches Out. pp. 74–81. Association for Computational Linguistics, Barcelona, Spain (Jul 2004), https://aclanthology.org/W04-1013/
15. Liu, Y., et al.: G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 2511–2522. Association for Computational Linguistics, Singapore (Dec 2023). https://doi.org/10.18653/v1/2023.emnlp-main.153
16. Luo, H., et al.: HyperGraphRAG: Retrieval-Augmented Generation via Hypergraph-Structured Knowledge Representation. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (Oct 2025), https://openreview.net/forum?id=ravS5h8MNg
17. Ostmeier, S., et al.: GREEN: Generative Radiology Report Evaluation and Error Notation. In: Findings of the Association for Computational Linguistics: EMNLP 2024. pp. 374–390. Association for Computational Linguistics, Miami, Florida, USA (Nov 2024). https://doi.org/10.18653/v1/2024.findings-emnlp.21
18. Papineni, K., et al.: Bleu: a Method for Automatic Evaluation of Machine Translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics. pp. 311–318. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA (Jul 2002). https://doi.org/10.3115/1073083.1073135
19. Sellergren, A., et al.: MedGemma Technical Report (Apr 2026). https://doi.org/10.48550/arXiv.2507.05201
20. Team, L., et al.: Lingshu: A Generalist Foundation Model for Unified Multimodal Medical Understanding and Reasoning (Jun 2025). https://doi.org/10.48550/arXiv.2506.07044
21. Wang, F., et al.: R2GenKG: Hierarchical Multi-modal Knowledge Graph for LLM-based Radiology Report Generation (Aug 2025). https://doi.org/10.48550/arXiv.2508.03426
22. Wang, X., et al.: MEDMKG: Benchmarking Medical Knowledge Exploitation with Multimodal Knowledge Graph (May 2025). https://doi.org/10.48550/arXiv.2505.17214
23. Wu, J.T., et al.: Chest ImaGenome Dataset for Clinical Reasoning. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Aug 2021), https://openreview.net/forum?id=H-d5634yVi
24. Yang, Y., et al.: HyperWalker: Dynamic Hypergraph-Based Deep Diagnosis for Multi-Hop Clinical Modeling across EHR and X-Ray in Medical VLMs (Jan 2026). https://doi.org/10.48550/arXiv.2601.13919
25. Yu, F., et al.: Evaluating progress in automatic chest X-ray radiology report generation. Patterns (Sep 2023). https://doi.org/10.1016/j.patter.2023.100802
26. Zhang, T., et al.: BERTScore: Evaluating Text Generation with BERT. In: International Conference on Learning Representations (Sep 2019), https://openreview.net/forum?id=SkeHuCVFDr
27. Zheng, L., et al.: Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In: Proceedings of the 37th International Conference on Neural Information Processing Systems. pp. 46595–46623. NIPS '23, Curran Associates Inc., Red Hook, NY, USA (Dec 2023), https://openreview.net/forum?id=uccHPGDlao