

ClinRAG-GRAPH: Clinical-prior Retrieval-Augmented Graph Model with Domain Adversarial Learning for Breast pCR Prediction

Yaofei Duan^{1,2,3}, Yuhao Huang^{4,5}, Tianyu Zhang^{1,2}, Yuan Gao^{2,6}, Luyi Han^{1,2}, Xin Wang^{2,6}, Xinyu Xie³, Xinglong Liang^{1,2}, Chunyao Lu^{1,2}, Muzhen He⁷, Patrick Pang³, Yue Sun³, Ning Mao⁸, Tao Tan³(✉), and Ritse Mann^{1,2}

¹ Department of Medical Imaging, Radboud University Medical Center, Nijmegen, The Netherlands

² The Netherlands Cancer Institute, Amsterdam, The Netherlands

³ Faculty of Applied Sciences, Macao Polytechnic University, Macau, China
taotan@mpu.edu.mo

⁴ Boston Children's Hospital, Harvard Medical School, Boston, USA

⁵ Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation, Chinese Academy of Sciences, Hong Kong, China

⁶ Imaging Division, University Medical Center Utrecht, Utrecht, The Netherlands

⁷ Department of Radiology, Fuzhou University Affiliated Provincial Hospital, Fuzhou, China

⁸ Department of Radiology, Yantai Yuhuangding Hospital, Shandong, China

Abstract. Neoadjuvant chemotherapy (NAC) response prediction is clinically important for treatment stratification in breast cancer. However, robust pre-treatment pathological complete response (pCR) prediction remains challenging due to insufficient cross-modal modeling, multicenter imaging heterogeneity, and weak evidence-grounded interpretability. We propose **ClinRAG-GRAPH**, a **Clinically informed Retrieval-Augmented Generation Graph** framework, for pre-treatment pCR prediction from DCE-MRI, structured clinical variables, and biopsy-derived pathological biomarkers. ClinRAG-GRAPH constructs an intra-patient clinical-prior graph and applies a prior-guided relation-aware graph convolutional network for structured multimodal representation learning. To improve cross-center robustness, we introduce a dual-branch domain-adversarial learning strategy to suppress protocol-related MRI bias while preserving pCR-relevant features. To enhance interpretability, we further incorporate large language model (LLM)-driven subgraph RAG module that retrieves clinically analogous historical cases and integrates retrieved evidence for pCR inference. We assemble a large-scale multicenter NAC breast cancer cohort for extensive validation, drawing from two public sources and three in-house centers. Results show that ClinRAG-GRAPH achieves AUCs of 0.815 on the internal test set and 0.774/0.712 on two external test sets, demonstrating robust pre-treatment pCR prediction across centers. The code is available at the anonymized link <https://github.com/miccai26-1181/ClinRAG-GRAPH>.

Keywords: Breast · Multimodal learning · Domain generalization.

1 Introduction

Neoadjuvant chemotherapy (NAC) is a standard treatment for patients with locally advanced or high-risk breast cancer, and pathological complete response (pCR) after NAC is a clinically important endpoint associated with favorable prognosis in several breast cancer subtypes [19,13]. However, pCR is determined only after surgery by pathological assessment. Therefore, accurate prediction of pCR from multimodal data may support individualized NAC treatment [9].

Recent work on NAC response prediction has evolved from imaging-only models to multimodal frameworks that combine MRI with clinical and pathological information [15,20]. MRI-based studies demonstrated the feasibility of imaging-driven prediction, while multicenter studies also revealed sensitivity to scanner and protocol heterogeneity, highlighting the risk of center-specific bias and limited cross-center generalization [8,22]. For pre-treatment pCR prediction, prior work has shown that integrating baseline MRI with clinical variables improves performance compared with unimodal settings [8]. Subsequent studies further suggested that adding clinical and pathological information can improve predictive accuracy and reinforce the value of multimodal baseline integration [16,12]. These results support multimodal modeling as a promising direction for pCR prediction. While some recent methods consider longitudinal or multi-time data for pCR prediction [4,13,21], we focus on the pre-treatment setting because it is most relevant for early treatment escalation for likely non-responders and de-escalation strategies for patients with high response probability.

However, several limitations remain for robust multicenter pre-treatment pCR prediction. First, many methods still rely on simple fusion operators (e.g., direct concatenation), which may not fully exploit clinically structured cross-modal dependencies. Related limitations have also been noted in recent multimodal studies [13,21]. Second, explicit domain-invariant learning to mitigate multicenter MRI heterogeneity is not consistently incorporated [22]. Third, most pCR models output only a scalar pCR probability without case-based evidence retrieval, which limits evidence-grounded interpretation in clinical workflows, motivating retrieval-augmented generation (RAG) inference strategies [2].

To address these limitations, we propose **ClinRAG-GRAPH**, a clinically informed retrieval-augmented graph framework for pre-treatment pCR prediction from multimodal data, including DCE-MRI, structured clinical variables, and biopsy-derived pathological biomarkers. The main contributions of our work are as follows. First, we encode clinically motivated prior knowledge into the graph construction and learn multimodal features with a relation-aware attention graph convolutional network (GCN). Second, we design adversarial learning strategy to mitigate multicenter MRI protocol heterogeneity and improve task-relevant representations. We further incorporate an LLM-based subgraph RAG module to support evidence-grounded and interpretable pCR inference. Third, we conduct a multicenter evaluation on a large-scale NAC breast cancer cohort comprising two public datasets and three in-house datasets, demonstrating competitive and consistent performance against five representative baselines.




Fig. 1: Overview of ClinRAG-GRAPH for breast pCR prediction. Emb., Tab., and Enc. denote embedding, table, and encoder; $\langle \text{mask} \rangle$ denotes the edge-prior mask, and scissors indicate gradient-reversal adversarial decoupling.

2 Method

We proposed ClinRAG-GRAPH for multicenter breast pCR prediction by integrating heterogeneous multimodal data into a unified intra-patient graph. For each patient i , we denote the MRI volume as X_i , the clinical/pathological attributes as t_i , the binary pCR label as $y_i \in \{0, 1\}$, and the acquisition center as $c_i \in \{1, \dots, C\}$. Our goal is to learn a predictor that is discriminative for y_i while being robust to domain shifts induced by c_i .

As shown in Fig. 1, ClinRAG-GRAPH first encodes the MRI volume using a DenseNet-121 encoder and constructs a clinical-prior graph G_i with edge-wise prior strengths p_{uv} . It then applies a prior-guided relational attentional R-GCN to obtain a patient embedding and predict pCR, while enforcing center-invariant imaging representations via a gradient-reversal domain classifier. We further employ an LLM-driven case-based subgraph RAG module to retrieve clinically analogous historical cases based on learned embeddings and attention signatures, and then fuse the neighbor-derived estimate with the R-GCN prediction for a more evidence-grounded and robust inference.

2.1 Hierarchical Semantic Clinical-prior Graph Construction

In multi-modal pCR prediction, purely data-driven graph learning faces two practical challenges: (i) homogeneous connectivity (e.g., fully connecting image-

table/table-table pairs) blurs heterogeneous multi-modal medical semantics, and (ii) enforcing priors via auxiliary losses constraints can destabilize optimization and introduce extra hyperparameters.

We therefore construct a hierarchical semantic clinical-prior graph at the patient level. For each patient i , we define an intra-patient directed graph $G_i = (V, E)$ with one imaging node and multiple clinical-/pathological- variable nodes. Each directed edge $(u, v) \in E$ is associated with a relation type $r(u, v)$ and a scalar prior weight $p_{uv} > 0$ that explicitly encodes the confidence of medical knowledge about this interaction. Guided by clinical practice guidelines and radiologist consensus [11,1,5,6], we organize edges into three semantic tiers:

- **Strong prior.** Guideline-level, stable associations (e.g., strong couplings among key molecular biomarkers), encoded with a fixed large $p_{uv} = 2$.
- **Soft prior.** Clinically agreed relations with inter-patient variability, encoded with a relatively large initialization of $p_{uv} = 0.8$.
- **Learnable prior.** Uncertain relations, weakly initialized (small $p_{uv} = 0.2$) and primarily determined by data (e.g., patient age with breast cancer subtypes).

To inject priors without auxiliary loss terms, we incorporate p_{uv} as a log-bias on attention logits only at the first graph layer (Sec. 2.2), providing early calibration toward clinically plausible message passing while keeping deeper layers fully learnable. We fix the temperature $t = 0.2$ in all experiments.

2.2 PRAttn-RGCN: Prior-guided Relational Attentional R-GCN

Follow R-GCN [18], given $G_i = (V, E)$, we perform relation-aware attentive message passing to fuse the imaging node with heterogeneous clinical/pathological variable nodes. Each directed edge $(u \rightarrow v)$ carries a semantic relation type $r(u, v)$, and nodes belong to different modalities. The prior strength p_{uv} is injected only at the first layer to guide early-stage cross-modal information flow.

Relation-aware attention. Let $\mathbf{h}_v^{(l)} \in \mathbb{R}^d$ be the node representation at layer l , and $\mathbf{t}_{\tau(v)}$ be the embedding of node type $\tau(v)$. For each edge $(u \rightarrow v)$, we compute a relation-specific attention logit

$$e_{uv}^{(l)} = \mathbf{a}_{r(u,v)}([\mathbf{h}_u^{(l)}; \mathbf{h}_v^{(l)}; \mathbf{t}_{\tau(u)}; \mathbf{t}_{\tau(v)}]), \quad (1)$$

where $\mathbf{a}_r(\cdot)$ is a lightweight relation-specific scorer. We inject the prior log-bias only when $l = 0$:

$$\tilde{e}_{uv}^{(l)} = e_{uv}^{(l)} + \mathbb{I}[l = 0] \cdot \tau \log(p_{uv}), \quad \alpha_{uv}^{(l)} = \text{softmax}_{u \in \mathcal{N}(v)}(\tilde{e}_{uv}^{(l)}). \quad (2)$$

Relational message passing. Messages are transformed per relation and aggregated with attention:

$$\mathbf{h}_v^{(l+1)} = \phi(\mathbf{W}_0 \mathbf{h}_v^{(l)} + \sum_{u \in \mathcal{N}(v)} \alpha_{uv}^{(l)} \mathbf{W}_{r(u,v)} \mathbf{h}_u^{(l)}), \quad (3)$$

where $\mathbf{W}_0$ means self-loop transform and $\mathbf{W}_r$ is relation-specific (optionally using basis decomposition with large $|\mathcal{R}|$). $\phi(\cdot)$ denotes LayerNorm&GELU.

Graph readout. After L layers, node embeddings are pooled to obtain a patient representation $\mathbf{g}_i = \text{MeanPool}(\{\mathbf{h}_v^{(L)}\}_{v \in V})$, which is fed to an Multilayer Perceptron (MLP) to predict pCR. The first-layer attention $\{\alpha_{uv}^{(0)}\}$ is retained as the attention signature for downstream case retrieval.

2.3 Adversarial Learning for Decoupling MRI Domain Noise

Multicenter DCE-MRI inevitably contains scanner/protocol-specific artifacts that can dominate image embeddings and harm cross-center generalization. To suppress such domain cues while preserving pCR-discriminative information, we impose a domain-adversarial constraint on the imaging representation. Let $\mathbf{h}_{\text{img}}^{(0)} = f_{\text{img}}(X_i) \in \mathbb{R}^d$ denotes the embedding used to initialize the imaging node. We attach an adversarial branch consisting of a lightweight projector $f_{\text{adv}}(\cdot)$ and a domain classifier $D(\cdot)$, trained through a gradient reversal layer (GRL) [3]:

$$\mathbf{z}_{\text{adv}} = f_{\text{adv}}(\mathbf{h}_{\text{img}}^{(0)}), \quad \hat{\mathbf{c}} = D(\text{GRL}(\mathbf{z}_{\text{adv}})), \quad (4)$$

where $\hat{\mathbf{c}} \in \mathbb{R}^C$ denotes the center logits and GRL multiplies the back-propagated gradient by a negative constant, enabling end-to-end optimization. Then, the training objective combines pCR supervision and domain confusion is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pCR}}(\hat{y}_i, y_i) + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}(\hat{\mathbf{c}}, c_i), \quad (5)$$

where $\mathcal{L}_{\text{pCR}}$ is BCE loss, $\mathcal{L}_{\text{adv}}$ is CE loss with λ_{adv} set to 0.05 through empirical tuning.

2.4 LLM-driven Subgraph RAG for Evidence-grounded Inference

To improve the robustness under distribution shift, we augment the trained PRAttn-RGCN with an LLM (**DeepSeek-V3.2** [10]) retriever, design an RAG module that grounds inference on similar patient evidence. After training, we infer the training set and export intermediate results as an evidence buffer. For each patient i , we store (i) a graph-level embedding $\mathbf{g}_i$, (ii) an MRI embedding $\mathbf{z}_i$, (iii) a clinical-pathological embedding $\mathbf{m}_i$, and (iv) edge-level attention weights α_i from the frozen PRAttn-RGCN. We then build FAISS [7] inner-product indices over ℓ_2 -normalized embeddings for approximate nearest-neighbor search.

At inference, a query patient q is encoded as $(\mathbf{g}_q, \mathbf{z}_q, \mathbf{m}_q, \alpha_q)$. DeepSeek outputs a schema-constrained retrieval plan from query-side signals (e.g., candidate pool sizes, filters (reduce false neighbors), fusion weights, and top- K), without making predictions. Candidates are retrieved via multi-route search and ranked by a fused score combining representation similarity and a subgraph match induced by α . The top- K neighbors are deterministically aggregated and combined with the PRAttn-RGCN output to yield the final prediction, enabling evidence-grounded inference while preserving backbone calibration. Optionally, DeepSeek generates a structured post-hoc explanation based solely on retrieval metadata, fused scores, and driver-edge cues, exported as a verifiable JSON record.

Table 1: Multicenter dataset summary across **internal** and **external** domains.

Domain		DUKE [17]	ISPY1 [14]	Zcenter	Ycenter	Qcenter
Amount	pCR	62	49	83	54	33
	non-pCR	235	123	367	166	88
Public available		Yes	Yes	No	No	No

Table 2: Performance comparison on internal and external test sets. AUC differences were evaluated using the paired DeLong test with a significance threshold of $p < 0.05$. Differences were significant at $p < 0.05$ for all comparisons.

Method	Internal (include 3 test sets)				External Ycenter				External Qcenter			
	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE
R-GCN	0.706 (0.611-0.792)	0.668	0.500	0.837	0.648 (0.560-0.730)	0.575	0.241	0.910	0.680 (0.580-0.778)	0.530	0.242	0.818
CLST	0.539 (0.450-0.626)	0.498	0.056	0.940	0.533 (0.412-0.651)	0.491	0.061	0.921	0.513 (0.409-0.614)	0.532	0.214	0.850
M2Fusion	0.505 (0.399-0.613)	0.534	0.167	0.902	0.484 (0.396-0.573)	0.455	0.037	0.873	0.454 (0.337-0.572)	0.491	0.061	0.921
LMF	0.678 (0.584-0.768)	0.625	0.452	0.797	0.693 (0.614-0.767)	0.510	0.074	0.946	0.666 (0.565-0.762)	0.532	0.212	0.852
iMRhpc	0.699 (0.605-0.790)	0.635	0.381	0.889	0.682 (0.598-0.761)	0.538	0.185	0.892	0.703 (0.605-0.792)	0.492	0.030	0.955
Ours	0.815 (0.738-0.885)	0.691	0.500	0.882	0.774 (0.700-0.843)	0.575	0.204	0.946	0.712 (0.604-0.815)	0.676	0.455	0.898

3 Experiments

Datasets and Implementation Details. We evaluated ClinRAG-GRAPH on two public datasets and three in-house datasets, with dataset statistics summarized in Table 1. Specifically, data from three source centers (DUKE [17], ISPY1 [14], and Zcenter) were used for model development and split at the patient level in a 7:1:2 ratio for training, validation, and internal testing, respectively, while Qcenter and Ycenter were held out for external generalization evaluation. All models were implemented in *PyTorch* and trained on an NVIDIA A40 GPU (48 GB memory). For preprocessing, multi-phase DCE-MRI scans were cropped to tumor-centered regions and concatenated along the temporal dimension to form an input volume of size (192, 256, 256), with intensities processed using z-score normalization. The network was optimized using AdamW (learning rate 2×10^{-5} , weight decay 1×10^{-4} , batch size 12) for 50 epochs.

Evaluation Metrics. The quantitative evaluation was performed using standard classification metrics, i.e., AUC (95% CI) to assess discrimination, balanced accuracy (bAcc) to account for class imbalance, sensitivity (SEN) to evaluate positive-case detection, and specificity (SPE) to assess negative-case recognition.

Internal Comparison. We evaluate ClinRAG-GRAPH against state-of-the-art methods including CLSTM [22], M2Fusion [21], LMF [13], iMRhpc [4]. As shown in Table 2, on the internal test set (aggregated from three internal test sets), ClinRAG-GRAPH achieves the best AUC of 0.815 (95% CI: 0.738–0.885),

Table 3: Ablation study results for model components, subgraph RAG is applied only to external sets. CP: clinical prior; PR: PRAttn-RGCN; ADV: adversarial decoupling; RAG: DeepSeek-based subgraph RAG. AUC differences were evaluated using the paired DeLong test ($p < 0.05$). Differences were significant at $p < 0.05$ except: V1 on Ycenter ($p = 0.074$), V2 on internal sets ($p = 0.062$).

Module	Internal (include 3 test sets)				External Ycenter				External Qcenter			
CPPRADVRAG	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE	AUC $\uparrow$ (95% CI)	bAcc $\uparrow$	SEN	SPE
- - - -	0.706 (0.611-0.792)	0.668	0.500	0.837	0.648 (0.560-0.730)	0.575	0.241	0.910	0.680 (0.580-0.778)	0.530	0.242	0.818
✓ - - -	0.745 (0.653-0.830)	0.679	0.548	0.811	0.708 (0.628-0.784)	0.606	0.278	0.934	0.683 (0.583-0.780)	0.610	0.515	0.705
✓ ✓ - -	0.787 (0.708-0.861)	0.645	0.452	0.837	0.746 (0.669-0.816)	0.603	0.296	0.910	0.709 (0.603-0.806)	0.697	0.576	0.818
✓ ✓ ✓ -	0.815 (0.738-0.885)	0.691	0.500	0.882	0.768 (0.695-0.835)	0.600	0.315	0.886	0.703 (0.593-0.804)	0.642	0.455	0.830
✓ ✓ ✓ ✓	-	-	-	-	0.774 (0.700-0.843)	0.575	0.204	0.946	0.712 (0.604-0.815)	0.676	0.455	0.898




Fig. 2: AUC results




Fig. 3: PLS-DA visualizations of binary class separation across internal and external cohorts.

outperforming R-GCN (0.706) and other fusion strategies such as LMF (0.678) and iMRhpc (0.699). It also improves balanced accuracy (0.691) and specificity (0.882) while maintaining sensitivity (0.500), indicating a more favorable trade-off. Paired DeLong tests confirm significant AUC gains ($p < 0.05$).

External evaluation. ClinRAG-GRAPH remains robust under cross-domain shift. On Ycenter, it attains the highest specificity (0.946) while preserving the best AUC, reducing false positives in unseen domains. On the more challenging Qcenter cohort, it yields the best balanced accuracy (0.676) and markedly higher sensitivity (0.455) than all baselines, reflecting improved detection without sacrificing overall discrimination. The PLS-DA projection in Fig. 3 further shows that model effectively distinguishes pCR patients across internal and external cohorts. Overall, ClinRAG-GRAPH consistently enhances both in-domain performance and multicenter generalization for pre-treatment pCR prediction.

Ablation Study. Table 3 shows that each proposed component contributes to performance. Introducing the clinical-prior graph improves over the R-GCN backbone, highlighting the value of explicit clinical-pathological relational priors for multimodal fusion. PRAttn-RGCN and adversarial decoupling provide further gains, suggesting the benefits of prior-aware message passing and more domain-invariant MRI representations. In addition, LLM-driven subgraph RAG




Fig. 4: SHAP analysis validating the contribution of key directed graph edges in the ClinRAG-GRAPH for LuminalA/B, HER2-pos, and TripleNeg subgroups.

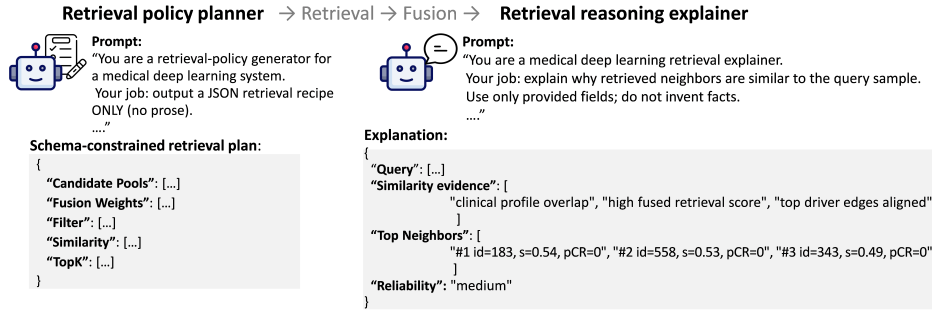


Fig. 5: LLM-driven RAG case study for a pCR=0 patient from external set.

further improves performance on external test sets, indicating enhanced robustness under distribution shift. Modality ablation (Fig. 2) further shows that pathology contributes most to pCR prediction.

Qualitative analysis. Fig. 4 presents subtype-specific SHAP analysis of directed edges in ClinRAG-GRAPH, showing that the model captures clinically meaningful cross-modal dependencies rather than relying on isolated features. The influential edges vary across Luminal A/B, HER2-pos, and TripleNeg subgroups, suggesting subtype-aware message passing. Consistent contributions from edges linked to ER, PR, HER2, Ki67, age, and MRI features support the effectiveness of our proposed hierarchical clinical-prior graph. The directional importance patterns further indicate that the model benefits from structured relational reasoning instead of simple feature concatenation.

LLM-driven RAG case study. As illustrated in Fig. 5, the Deepseek serves as both a retrieval planner and a post-hoc explainer. The planner generates a schema-constrained retrieval policy that selects clinically matched neighbors for the pCR=0 query case. The retrieved samples show high fused similarity and consistent clinical profiles. After fusion, the top-3 neighbors provide coherent supporting evidence with aligned driver-edge patterns. The LLM explainer then

summarizes this evidence into a structured, verifiable rationale. The final prediction agrees with the retrieved evidence and remains stable for this case.

4 Conclusion

In this study, we explore pre-treatment pCR prediction under multimodal heterogeneity and multicenter domain shift. We propose ClinRAG-GRAPH, a clinically informed framework that integrates DCE-MRI, clinical variables, and pathological biomarkers through hierarchical graph modeling, adversarial learning, and RAG inference. Extensive experiments demonstrate improved predictive performance, cross-center generalization, and clinically interpretability. Future work will collect longitudinal data to further improve pCR prediction.

Acknowledgments. This work was supported by Shenzhen Medical Research Fund (D2501013), Macao Polytechnic University Grant (RP/FCA-13/2026), and the Science and Technology Development Fund, Macau SAR (File no. 0004/2025/ASJ) under the FDCT-FAPESP Joint Funding Scheme.

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

References

1. Amin, M.B., Edge, S.B., Greene, F.L., Byrd, D.R., Brookland, R.K., Washington, M.K., Gershenwald, J.E., Compton, C.C., Hess, K.R., Sullivan, D.C., et al.: AJCC cancer staging manual, vol. 1024. Springer (2017)
2. Amugongo, L.M., Mascheroni, P., Brooks, S., Doering, S., Seidel, J.: Retrieval augmented generation for large language models in healthcare: A systematic review. *PLOS Digital Health* **4**(6), e0000877 (2025)
3. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of machine learning research* **17**(59), 1–35 (2016)
4. Gao, Y., Ventura-Diaz, S., Wang, X., He, M., Xu, Z., Weir, A., Zhou, H.Y., Zhang, T., van Duijnhoven, F.H., Han, L., et al.: An explainable longitudinal multi-modal fusion model for predicting neoadjuvant therapy response in women with breast cancer. *Nature communications* **15**(1), 9613 (2024)
5. Gradishar, W.J., Moran, M.S., Abraham, J., Aft, R., Agnese, D., Allison, K.H., Anderson, B., Burstein, H.J., Chew, H., Dang, C., et al.: Breast cancer, version 3.2022, nccn clinical practice guidelines in oncology. *Journal of the National Comprehensive Cancer Network* **20**(6), 691–722 (2022)
6. Hurvitz, S.A., Martin, M., Symmans, W.F., Jung, K.H., Huang, C.S., Thompson, A.M., Harbeck, N., Valero, V., Stroyakovskiy, D., Wildiers, H., et al.: Neoadjuvant trastuzumab, pertuzumab, and chemotherapy versus trastuzumab emtansine plus pertuzumab in patients with her2-positive breast cancer (kristine): a randomised, open-label, multicentre, phase 3 trial. *The Lancet Oncology* **19**(1), 115–126 (2018)
7. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *IEEE transactions on big data* **7**(3), 535–547 (2019)

8. Joo, S., Ko, E.S., Kwon, S., Jeon, E., Jung, H., Kim, J.Y., Chung, M.J., Im, Y.H.: Multimodal deep learning models for the prediction of pathologic response to neoadjuvant chemotherapy in breast cancer. *Scientific reports* **11**(1), 18800 (2021)
9. Krasniqi, E., Filomeno, L., Arcuri, T., Ferretti, G., Gasparro, S., Fulvi, A., Roselli, A., D'Onofrio, L., Pizzuti, L., Barba, M., et al.: Multimodal deep learning for predicting neoadjuvant treatment outcomes in breast cancer: a systematic review. *Biology Direct* **20**(1), 72 (2025)
10. Liu, A., Feng, B., Xue, B., Wang, B., Wu, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., et al.: Deepseek-v3 technical report. arXiv preprint arXiv:2412.19437 (2024)
11. Loibl, S., André, F., Bachelot, T., Barrios, C., Bergh, J., Burstein, H., Cardoso, M., Carey, L., Dawood, S., Del Mastro, L., et al.: Early breast cancer: Esmo clinical practice guideline for diagnosis, treatment and follow-up. *Annals of Oncology* **35**(2), 159–182 (2024)
12. Lv, M., Zhao, B., Mao, Y., Wang, Y., Su, X., Zhang, Z., Wu, J., Gao, X., Wang, Q.: Deep learning model for the early prediction of pathologic response following neoadjuvant chemotherapy in breast cancer patients using dynamic contrast-enhanced mri. *Frontiers in Oncology* **15**, 1491843 (2025)
13. Ma, D., Cao, J., Cheng, H., Zhou, D., Liu, J., Zhang, X., Wu, K., Gu, C., Guan, X.: Longitudinal mri-clinical multimodal fusion for pcr prediction in breast cancer. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 323–332. Springer (2025)
14. Newitt, D.C., Hylton, N.M., on behalf of the I-SPY1 Network, Team, A.T.: Multi-center breast dce-mri data and segmentations from patients in the i-spy 1/acrin 6657 trials [data set] (2016). <https://doi.org/10.7937/K9/TCIA.2016.HdHpgJLK>, <https://doi.org/10.7937/K9/TCIA.2016.HdHpgJLK>
15. Nishizawa, T., Maldjian, T., Jiao, Z., Duong, T.Q.: Attention-based multimodal deep learning for interpretable and generalizable prediction of pathological complete response in breast cancer. *Journal of Translational Medicine* **23**(1), 774 (2025)
16. Pesapane, F., Rotili, A., Botta, F., Raimondi, S., Bianchini, L., Corso, F., Ferrari, F., Penco, S., Nicosia, L., Bozzini, A., et al.: Radiomics of mri for the prediction of the pathological response to neoadjuvant chemotherapy in breast cancer patients: a single referral centre analysis. *Cancers* **13**(17), 4271 (2021)
17. Saha, A., Harowicz, M.R., Grimm, L.J., Weng, J., Cain, E.H., Kim, C.E., Ghate, S.V., Walsh, R., Mazurowski, M.A.: Dynamic contrast-enhanced magnetic resonance images of breast cancer patients with tumor locations (duke-breast-cancer-mri) [data set] (2021). <https://doi.org/10.7937/TCIA.e3sv-re93>, <https://doi.org/10.7937/TCIA.e3sv-re93>
18. Schlichtkrull, M., Kipf, T.N., Bloem, P., Van Den Berg, R., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: *European semantic web conference*. pp. 593–607. Springer (2018)
19. Spring, L.M., Fell, G., Arfe, A., Sharma, C., Greenup, R., Reynolds, K.L., Smith, B.L., Alexander, B., Moy, B., Isakoff, S.J., et al.: Pathologic complete response after neoadjuvant chemotherapy and impact on breast cancer recurrence and survival: a comprehensive meta-analysis. *Clinical cancer research* **26**(12), 2838–2848 (2020)
20. Xu, Z., Zhou, Z., Son, J.B., Feng, H., Adrada, B.E., Moseley, T.W., Candelaria, R.P., Guirguis, M.S., Patel, M.M., Whitman, G.J., et al.: Deep learning models based on pretreatment mri and clinicopathological data to predict responses to neoadjuvant systemic therapy in triple-negative breast cancer. *Cancers* **17**(6), 966 (2025)

21. Zhang, S., Du, S., Sun, C., Li, B., Shao, L., Zhang, L., Wang, K., Liu, Z., Tian, J.: M2fusion: multi-time multimodal fusion for prediction of pathological complete response in breast cancer. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 458–468. Springer (2024)
22. Zhang, Y., Chen, J.H., Lin, Y., Chan, S., Zhou, J., Chow, D., Chang, P., Kwong, T., Yeh, D.C., Wang, X., et al.: Prediction of breast cancer molecular subtypes on dce-mri using convolutional neural network with transfer learning between two centers. *European radiology* **31**(4), 2559–2567 (2021)