BreastGPT: A Multimodal Large Language Model for the Full Spectrum of Breast Cancer Clinical Routine

Yang Liu[1,*], Jiajin Zhang[1,3,*], Danyang Tu[1,3,†,*], Yaojun Hu[2], Jiao Qu[4,5], Jiuyu Zhang[6], Yu Shi[6], Wei Fang[1,3,*], Shi Gu[2,†], Ling Zhang[1], Yingda Xia[1]

[1]DAMO Academy, Alibaba Group [2]Zhejiang University
[3]Hupan Lab [4]Department of Radiology, West China Hospital, Sichuan University
[5]Institute of Breast Health Medicine, West China Hospital, Sichuan University
[6]China Medical University

[*]Equal contribution. [†]Corresponding authors.

Abstract

Breast cancer care requires multimodal reasoning across a clinical workflow spanning *screening*, *diagnosis*, and *treatment planning*. However, existing medical MLLMs are typically limited to isolated modalities or narrow task settings, restricting their ability to support workflow-level clinical reasoning. To address this gap, we introduce **BreastStage**, a workflow-aligned breast imaging instruction corpus containing 1.86M instruction-following pairs from 17 sub-datasets, 5 imaging modalities, and 136 task templates, together with **BreastStage-Bench**, a benchmark for evaluating multimodal reasoning across the breast cancer care continuum. Built upon this corpus, we propose **BreastGPT**, a unified MLLM with a dual-branch visual encoder and concept-preserving token compression to bridge standard radiology and gigapixel pathology. On BreastStage-Bench, BreastGPT achieves 75.66% closed-ended accuracy and 89.92% open-ended score, outperforming both general-purpose and medical-specific MLLMs across clinical stages and task formats. Data, code, and model checkpoints are publicly released at `https://yangyy-liu.github.io/BreastGPT.io`.

1 Introduction

Breast cancer remains one of the most prevalent and life-threatening malignancies among women worldwide (Arnold et al., 2022; Freihat et al., 2025; Sha et al., 2024; Kim et al., 2025). Its clinical management involves not just a single diagnostic decision, but a staged workflow spanning *screening*, *diagnosis*, and *treatment planning*. Each stage relies on distinct imaging evidence and pursues different clinical objectives: *Screening* prioritizes triage at the population level, where mammography, breast ultrasound (BUS) and chest CT are mainly used to detect early suspicious findings and estimate their risk (Katsika et al., 2024); *Diagnosis* then shifts toward fine-grained lesion characterization, where multimodal MRI is pivotal for evaluating morphology and probability of malignancy (Aristokli et al., 2022); *Treatment planning* requires tumor biology and response assessment, for which gigapixel whole-slide images (WSIs) provide pathological subtyping and biomarker evidence (McCaffrey et al., 2024). This clinical continuum imposes a different requirement from conventional single-task image understanding: a clinically useful model must process heterogeneous modalities while preserving the stage-specific reasoning style expected at each point of care.

However, despite rapid progress in medical MLLMs, current breast imaging AI remains fragmented: datasets, models, and benchmarks are invariably organized around isolated imaging modalities and limited tasks for a single clinical stage. Consequently, 1) these single-modality datasets fail to comprehensively evaluate models' performance across the various stages within the practical clinical workflow; 2) there is a lack of a versatile model capable of processing different modalities uniformly and efficiently to address various tasks at different stages of real-world clinical workflows.

To this end, we first construct BreastStage (Fig. 1), a large-scale multimodal instruction corpus aligned with the breast cancer clinical workflow. BreastStage integrates **17** sub-datasets across **5** modalities (mammography, BUS, MRI, WSI, and CT), yielding **1.86M** instruction-following QA pairs over **136** task templates. These instructions cover attribute extraction, report generation, visual question answering, and image-grounded captioning, providing supervision over both radiological and pathological evidence. Based on the held-out split of BreastStage, we further build **BreastStage-Bench**, a comprehensive benchmark for workflow-level breast oncology reasoning developed with breast oncology experts. We assess multiple publicly available multimodal large language mod-

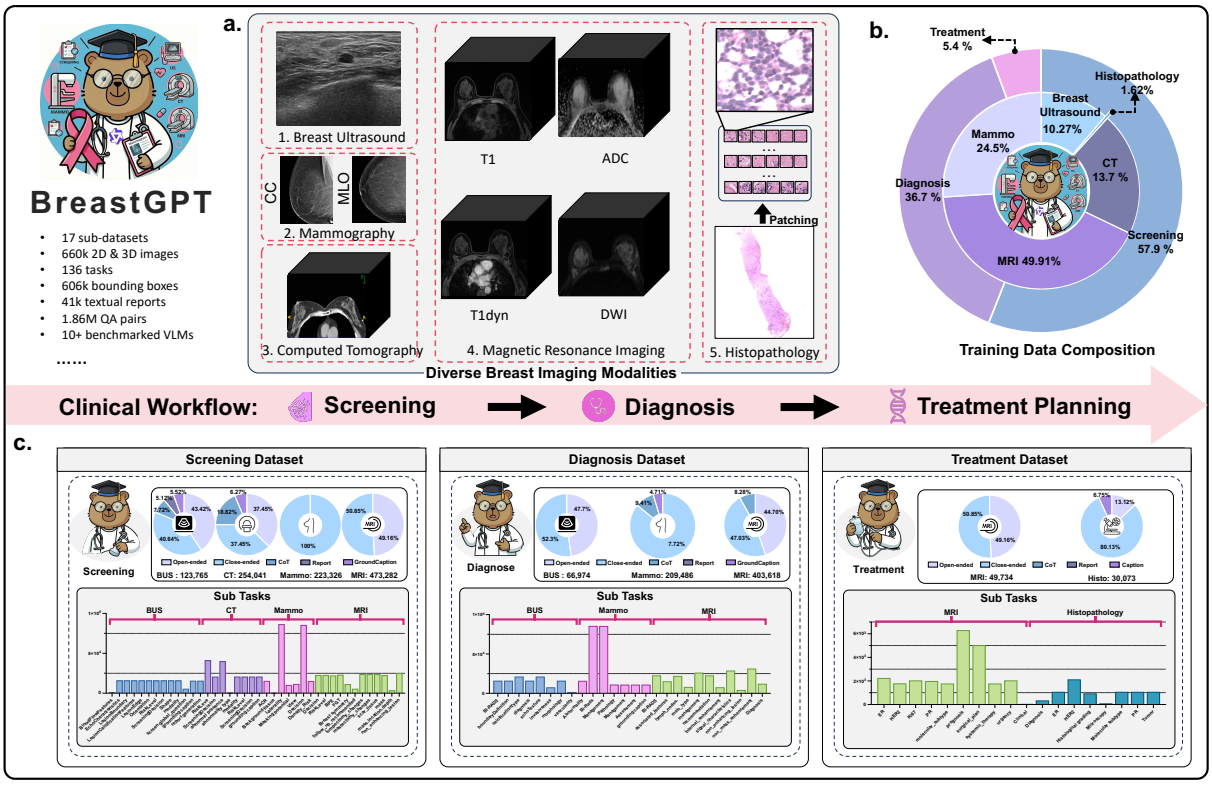


Figure 1: Overview of the **BreastStage** dataset aligned with the end-to-end breast cancer clinical workflow. The dataset is structured across three clinical stages: (1) **Screening** utilizes mammography and breast ultrasound (BUS) images for early lesion detection; (2) **Diagnosis** leverages MRI (main modality), BUS and Mammography for detailed tumor characterization; and (3) **Treatment planning** involves gigapixel whole-slide images (WSI) for pathological subtyping and therapeutic evaluation.

els, including both general-purpose and medical-specific variants, alongside advanced proprietary models. Extensive experiments demonstrate that BreastStage-Bench poses a significant challenge for current MLLMs: GPT-5.4 achieves an average score of only 54%, and existing medical MLLMs show no clear advantage over general-purpose models on these clinical workflow-aligned tasks.

Based on BrasetStage, we propose **BreastGPT**, a unified MLLM for breast cancer workflow reasoning. BreastGPT handles all five imaging modalities within a single architecture and adopts stage-conditioned system prompts to switch between screening, diagnostic, and treatment-oriented reasoning behaviors. Meanwhile, to address the significant differences in image scale across modalities, BreastGPT adopts a dual-branch visual encoder with a resolution-aware gating module that automatically routes each input to an appropriate processing branch based on stage condition. Aligned with practical clinical workflow, BreastGPT demonstrates exceptional versatility and achieves remarkable performance. Compared to existing general-purpose and medical-specific MLLMs, BreastGPT achieves performance gains of over 25%, 35%, and 40% for screening, diagnosis, and treatment planning, respectively. To summarize, our contributions are threefold:

Workflow-level problem formulation and benchmark. We formulate breast cancer multimodal reasoning as a workflow-aligned problem spanning screening, diagnosis, and treatment planning. Based on this, we construct BreastStage, a large-scale instruction corpus with 1.86M QA pairs from 17 sub-datasets across 5 modalities, and BreastStage-Bench, a held-out benchmark for evaluating stage-specific clinical reasoning.

Unified cross-scale breast MLLM. We introduce BreastGPT, a unified multimodal LLM that handles mammography, ultrasound, CT, MRI, and gigapixel pathology within a single instruction-following framework, enabling consistent reasoning across the full breast cancer clinical routine.

Concept-preserving visual token compression. To bridge the extreme scale gap between standard radiology and WSIs, we propose a dual-branch visual encoder with a concept-based token selector that preserves clinically salient visual evidence un-

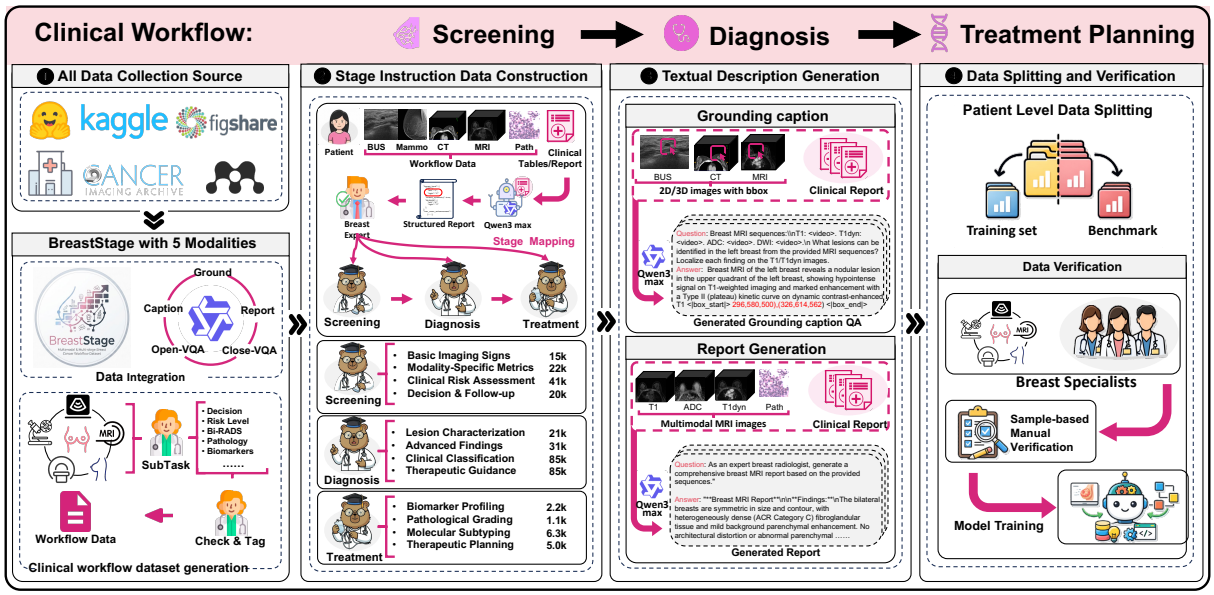


Figure 2: **BreastStage construction pipeline.** It involves curating diverse imaging modalities from public datasets and hospitals. The raw data is then cleaned, annotated, and checked by breast experts. Textual and instruction data are generated by prompting Qwen3-Max. Finally, the data is split into a training set and a benchmark, with professional breast experts assessing the training samples.

der a fixed token budget, substantially improving both accuracy and inference efficiency.

2 Related Work

Medical MLLMs. Medical MLLMs such as LLaVA-Med (Lu et al., 2024), HuatuoGPT-Vision (Chen et al., 2024a), and HealthGPT (Lin et al., 2025b), together with a growing family of generalist medical vision–language systems (Chen et al., 2023; Xu et al., 2025; Xiao et al., 2025; Ye and Tang, 2025; Kalpélbé et al., 2025; AlSaad et al., 2024), aim to unify diverse medical tasks under one instruction-following interface. Broad medical benchmarks (Hu et al., 2024; Chen et al., 2024b; Yue et al., 2024; Liu et al., 2024; Duan et al., 2024; Hu et al., 2026) further reveal that strong general vision–language ability does not directly yield reliable cross-specialty clinical reasoning. These efforts emphasize cross-specialty breadth rather than the *workflow structure* of a single disease.

Breast Imaging AI. Existing breast imaging models are largely organized around a single modality: VersaMammo (Huang et al., 2025) and MammoVQA (Zhu et al., 2025) for mammography, BUS-CoT (Yu et al., 2026) for ultrasound, and MOME (Luo et al., 2025) for MRI. Despite strong modality-specific performance, none reasons jointly across the screening-diagnosis-treatment pathway. BreastGPT unifies five imaging modalities within a single framework and employs a dual-branch encoder to bridge the scale gap between standard radiology and gigapixel pathology. Additional related work is provided in § B.

3 BreastStage Dataset Curation

3.1 Scope and scale

BreastStage contains about 662,000 unique 2D and 3D images, 136 task templates, 606,226 records with bounding-box or mask annotations, and 1.86 million instruction-following pairs. Task coverage spans classification, visual grounding, open- and closed-ended question answering, captioning, and report generation, so the dataset measures both perception accuracy on narrow subtasks and the model's ability to follow diverse clinical instructions. Samples are distributed as 57.9% screening, 36.7% diagnosis, and 5.4% treatment, broadly tracking the relative prevalence and data availability of these stages in real breast care. Please check § C.1 for details.

3.2 Data construction pipeline

As illustrated in Fig. 2, the BreastStage curation pipeline has four steps:

i. Workflow Data Generation. We assemble a cohort of breast imaging studies that draws CT, BUS, mammography, and WSI from public repositories and adds multimodal MRI from two collaborating clinical institutions. For quality control, modality-specific visual specialist agents built on

Qwen2.5-VL-72B (Bai et al., 2025b) emit a structured `{validity, reason}` verdict for each image; samples flagged as low-quality are routed to a Low-Image Check in which a breast specialist confirms or restores the verdict, rather than being silently dropped. Lesion-level spatial annotations are then attached: BUS reuses the expert hand-drawn masks of BUS-CoT (Yu et al., 2026), mammography reuses the EMBED-released bounding boxes, CT volumes from CT-RATE (Hamamci et al., 2026) are passed through DRT-M3D (Zhou et al., 2026) to obtain segmentations, and MRI carries manual T1 and dynamic contrast-enhanced T1dyn annotations from 10 breast specialists. The resulting cohort is tagged with one of 136 specialist-designed sub-tasks (`task template`) and audited by hand. Full data provenance and per-modality counts are reported in Table 5 and the complete pipeline is detailed in § C.2.

ii. Stage Instruction Data Construction. All data is standardized, cleaned, and paired with its matching clinical tables or reports, then converted into instruction-following QA pairs: *1)* Chinese reports are first translated to English by Qwen3-Max acting as a radiologist agent; *2)* Each translated report is then parsed into a modality-specific and BI-RADS-aligned structured report whose schema is designed by breast experts following clinical reporting guidelines; *3)* For every `task template` from the previous stage, breast experts then write closed-ended questions whose answer is drawn from one or more leaf fields of the structured report, with options taken directly from the schema's enum values, so the resulting QA pairs cannot introduce facts beyond the structured record; *4)* Finally, each `task template` is mapped, with expert guidance, to one of the three clinical stages: screening, diagnosis, or treatment.

After the stage mapping, the resulting BreastStage instruction set is organized to match the cognitive demands of each clinical stage, with category-level counts summarized in Table 6. Screening tasks emphasize early detection and triage, covering modality / view / laterality recognition, breast-tissue characterization, (FGT, BPE, density), risk and decision tasks, lesion presence and morphology, and post-surgical change tracking, totaling 1,075,092 pairs. Diagnosis tasks emphasize tumour characterization and staging, including BI-RADS, pathology classification, lesion-level signal and kinetic characterization, and associated findings such as lymph-node involvement and tissue invasion, totaling 680,409 pairs. Treatment tasks are grounded in histopathology and treatment-relevant radiology, including prognosis prediction, surgical planning, systemic therapy and biomarker assessment, and WSI-based pathology subtyping, totaling 100,567 pairs. The hierarchical organization ensures that instruction tuning covers the distinct reasoning patterns required across the patient journey. Please check § C.2 for details.

iii. Textual Description Generation. The same structured records that feed VQA construction are also reused to synthesise two narrative task types: ground captions and report generation: A *ground caption* encodes 2D or 3D spatial coordinates of lesions together with their categories and anatomical locations, providing a multi-dimensional spatial observation for BUS, mammography, CT, and MRI. We construct ground captions by prompting Qwen3-Max to integrate the original clinical report with the available bounding-box annotations into a single visually grounded narrative. *Medical reports*, in contrast, condense the multimodal evidence into key findings, tissue characteristics, and diagnostic or therapeutic assessments. After consultation with senior breast oncologists, we organize each generated report into the four ACR-aligned sections *Findings*, *Impression*, *Final Assessment*, and *Management*. The report generation pipeline is iterative: Qwen3-Max first drafts a comprehensive report from the multimodal input, including multiparametric MRI sequences and histopathology WSIs; breast specialists then audit a sample, catalog the recurring reasoning errors and hallucinations, and we refine the modality-specific system prompts before regenerating.

iv. Data Splitting and Verification. To prevent data leakage while preserving task diversity, we enforce strict patient-level separation between the training set and the evaluation benchmark. We allow the same image to be reused across different task contexts (for instance, a BUS image may carry both a BI-RADS label and a lesion-characterization task), while every record derived from a given patient is kept on a single side of the split. The split itself uses stratified sampling on the composite key of modality, task type, and pathology label, so the underlying task and class distributions stay consistent across train and test. Downstream evaluation therefore reflects clinical generalization rather than memorization of shared visual content.

For data reliability we run a multi-stage verification pipeline. Automated heuristic filters first

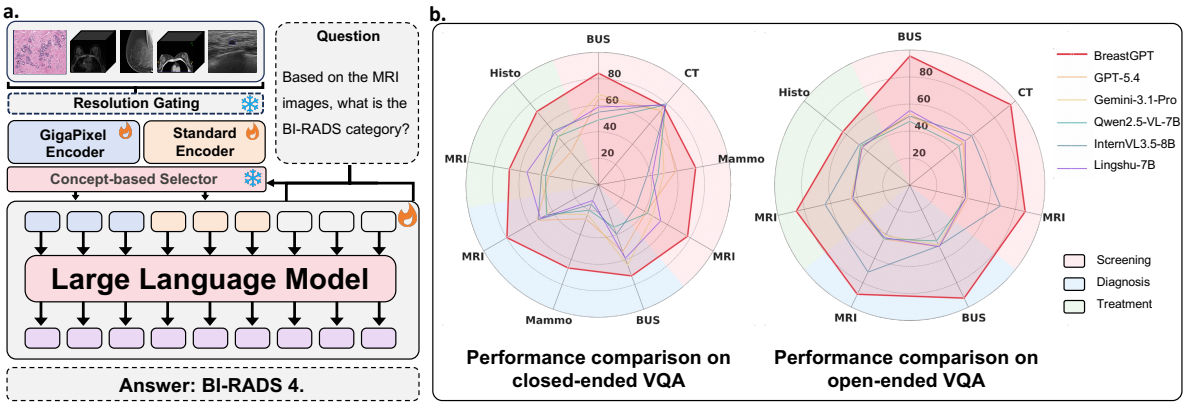


Figure 3: (a) **Architecture of BreastGPT.** A resolution gating mechanism routes standard radiological images to a ViT-based **Standard Branch**, and extreme-resolution WSIs to a specialized **GigaPixel Branch**. The WSI pipeline extracts patch features via a frozen CONCHv1.5 encoder, aggregates global context using LongNet, and compresses the sequence via a universal concept-based token selector before LLM injection. (b) **Performance Comparison.** BreastGPT demonstrates superior reasoning ability across both clinical QA tasks.

remove hallucinated clinical terms, malformed bounding-box coordinates, instruction-answer conflicts, and near-duplicate QA pairs detected by MinHash similarity above 0.85. Three breast specialists then conduct an independent task-level audit on the held-out test partition (5 random samples per task per specialist), and any task with consistent flags triggers refinement of the generation prompts or filtering rules. **Notably**, the clinical reliability of all generated data has been audited by clinical experts. In § E.3, we report expert agreement on a stratified subset of BreastStage-Bench, including task validity, answer correctness, and clinical consistency. The expert audit covers all modalities and task families, and disagreements are resolved by consensus review.

4 BreastGPT

BreastGPT is initialized from Qwen3-VL (Bai et al., 2025a), which natively supports multi-image inputs and inherits the resolution-aware visual modeling of the Qwen-VL lineage (Bai et al., 2023; Wang et al., 2024; Bai et al., 2025b). We keep the LLM backbone and adapt it to breast cancer reasoning through stage-aware system prompts that act as lightweight task routers at inference, so one model serves the screening → diagnosis → treatment workflow across 136 task templates without task-specific heads. To handle scales from radiology to gigapixel pathology, a *modality-aware resolution gating* module routes WSI inputs to the GigaPixel branch and CT, MRI, BUS, and mammography to the Standard branch (Fig. 3a).

4.1 Dual-Branch Visual Encoding

Standard Encoder for Radiology. Radiological modalities (CT, MRI, BUS, mammography) are encoded by the native Vision Transformer (ViT) of the Qwen3-VL backbone. This branch is optimized to capture macroscopic anatomical structures, tissue densities, and localized lesion textures that are critical for screening and diagnostic reasoning.

GigaPixel Encoder for Pathology. To integrate extreme-resolution histopathology without overwhelming the LLM context window, we introduce a specialized WSI pathology branch. Prior work on pathology foundation models and WSI transformers has shown that both spatially preserved local morphology and slide-level context are essential for reliable gigapixel image understanding (Ahmed et al., 2024; Lu et al., 2024; Chen et al., 2023, 2024c; Guo et al., 2025; Buzzard et al., 2024; Li et al., 2024; Xiong et al., 2025). Pathology information must therefore be incorporated without collapsing tens of thousands of patch-level features into clinically uninformative representations. We separate WSI processing into two steps:

- *Feature Extraction:* WSI tiles at 20× magnification are encoded by a frozen CONCH v1.5 foundation model (Lu et al., 2024), producing 512-dimensional patch embeddings $\{p_1, \ldots, p_N\}$, where N can reach tens of thousands. Using a frozen encoder retains robust pathology representations from large-scale pretraining while avoiding the computational instability of retraining.

- *Context Aggregation:* A trainable LongNet dilated attention encoder (Ding et al., 2023)

aggregates local cellular details and global architectural patterns across the full patch sequence, producing contextualized representations $\{h_1, \ldots, h_N\}$. Because diagnostic cues arise from both local cell morphology and broader tissue organization, LongNet's exponentially growing dilation rates preserve long-range context without the $O(N^2)$ bottleneck of standard attention.

Full implementation details (hyperparameters, dilation schedule, projection dimensions) are in § D.2.

4.2 Universal Concept-Based Compression

A unified model spanning 3D radiological volumes and gigapixel pathology faces a severe token-budget mismatch, as a single WSI patch sequence routinely exceeds tens of thousands of tokens. Rather than resorting to modality-specific pooling or heuristic truncation, we adapt the MM-Tok multimodal coverage maximization framework (Dong et al., 2026) as a universal concept-based token selector, following recent work on efficient VLM token pruning and compact visual representations (Shang et al., 2025; Tang and Shen, 2025; Sun et al., 2025; Meng et al., 2025; Omri et al., 2025). Given either redundant radiological ViT tokens or contextualized WSI embeddings $\mathbf{X} = \{x_i\}_{i=1}^{N}$, the selector distills the input into exactly k tokens by maximizing a concept-coverage objective:

$$\max_{S \subseteq \{1,\ldots,N\}, |S|=k} \underbrace{\frac{1}{M} \sum_{j=1}^{M} \max_{i \in S} \operatorname{sim}(t_j, x_i)}_{\text{text–vision coverage}} + \alpha \underbrace{\frac{1}{N} \sum_{j=1}^{N} \max_{i \in S} \operatorname{sim}(x_j, x_i)}_{\text{vision–vision coverage}} \quad (1)$$

where $\{t_j\}_{j=1}^{m}$ are text tokens from the clinical instruction and α balances query relevance with visual representativeness. The first term anchors selection to prompt-relevant clinical concepts, while the second term prevents collapse onto redundant salient regions and encourages coverage of global tissue or anatomical structure. This retains visually distinct, clinically informative patterns, enabling BreastGPT to condense multi-scale visual evidence into a fixed token budget without sacrificing the diversity required for downstream reasoning. The complete calibration and greedy optimization procedure are provided in § D.3.

5 Experiments

5.1 Experimental Setups

Baselines. We evaluate BreastGPT against three baseline groups on BreastStage-Bench, all run zero-shot with the same stage-specific prompts: five proprietary frontier models (GPT-5.4, Claude-opus-4-6, Claude-sonnet-4-6, Gemini-3.1-Flash, Gemini-3.1-Pro (Hurst et al., 2024; Team et al., 2023)); six open-source general-purpose VLMs at comparable scale (Qwen2.5-VL-3B/7B (Bai et al., 2025b), Qwen3-VL-4B/8B (Bai et al., 2025a), MiMo-VL-SFT (Team et al., 2025a), InternVL3.5); and two medical VLMs (Lingshu (Xu et al., 2025), HuatuoGPT-V (Chen et al., 2024a)), with broader medical-VLM context from (Li et al., 2023; He et al., 2024; Lin et al., 2025b). Per-baseline configuration is in § E.4.

BreastGPT variants. To isolate the effect of token selection, we run two variants that share backbone, branches, and training data and differ only in how the final 128 tokens are obtained. BreastGPT (cluster) uses our training-free greedy coverage selector; BreastGPT (learn) swaps it for a learnable cross-attention selector in the spirit of Flamingo's Perceiver Resampler (Alayrac et al., 2022).

5.2 Evaluation Metrics

Closed-ended tasks are scored with accuracy. For open-ended VQA and report generation we report BERTScore F1, BLEU, and ROUGE-1, plus a weighted aggregate ($0.5{\cdot}$BERTScore $+\,0.25{\cdot}$BLEU $+\,0.25{\cdot}$ROUGE-1); BERTScore takes the larger weight because clinical free-text varies in surface form even when the underlying finding is the same. Grounding uses mean IoU, with a sample credited IoU $= 1$ when the model and ground truth both abstain (TN) and 0 when the model hallucinates a box that does not exist (FP). See § E.2.

5.3 Evaluation Results

Overall observations. BreastGPT is the top scorer in every cell of Sections 5 and 5.3. The cleanest picture comes from closed-ended VQA, where the cluster variant averages 75.66 while every baseline group sits in the low fifties or below. Open-ended VQA widens the gap further: no main-cohort baseline clears 55, and even the strongest models

		Closed-ended VQA (Accuracy, %)										Open-ended VQA (normalized Score, %)							
		Screening				Diagnosis			Treatment			Screening			Diagnosis		Treatment		
Model	#P	BUS	CT	Mam	MRI	BUS	Mam	MRI	MRI	His	Avg	BUS	CT	MRI	BUS	MRI	MRI	His	Avg
Proprietary Models																			
GPT-5.4	–	64.89	78.55	68.51	41.43	53.46	55.26	53.50	38.10	32.28	54.00	53.43	49.34	59.85	50.42	59.20	58.07	44.73	53.58
Claude-opus-4-6	–	50.21	72.00	39.57	50.48	38.83	7.66	45.10	41.27	25.94	41.23	43.32	42.77	43.75	42.80	44.42	42.91	40.84	42.97
Claude-sonnet-4-6	–	65.53	68.00	37.02	40.00	54.79	19.67	55.94	42.86	36.30	46.68	43.63	43.12	44.31	43.00	44.98	42.62	41.06	43.25
Gemini-3.1-Flash	–	55.11	42.67	49.68	38.57	37.60	24.32	31.82	48.41	35.81	40.44	48.76	48.21	47.38	45.80	47.53	44.15	42.81	46.38
Gemini-3.1-Pro	–	68.09	73.33	50.21	47.14	64.36	23.87	43.88	44.44	46.53	51.32	51.19	51.46	41.60	48.94	41.39	42.80	45.77	46.16
Open-Source Models																			
Qwen2.5-VL	3B	46.81	51.39	33.51	45.71	30.85	12.76	51.75	44.44	51.52	40.97	49.52	48.39	49.36	47.36	50.54	47.04	44.80	48.14
Qwen2.5-VL	7B	49.15	**79.27**	44.47	44.76	34.31	14.41	46.85	37.30	47.62	44.24	46.20	46.86	49.47	44.99	49.67	44.70	43.98	46.55
Qwen3-VL	4B	52.13	75.39	41.91	49.05	35.37	21.17	48.08	43.65	38.61	45.04	45.20	55.49	47.71	44.93	51.46	48.98	41.26	47.86
Qwen3-VL	8B	57.87	78.55	39.68	51.43	48.14	25.83	47.90	34.92	47.02	47.93	44.94	48.53	44.85	44.34	46.24	44.15	41.21	44.89
MiMo-VL	7B	15.11	23.52	8.62	6.67	14.89	6.16	0.35	0.79	19.98	10.68	42.49	55.98	63.70	41.69	65.66	63.08	39.86	53.21
InternVL3.5	8B	51.91	77.70	35.85	34.76	42.55	16.22	52.27	44.44	52.98	45.41	50.74	58.05	56.56	48.80	58.89	55.94	46.50	53.64
Medical-Specific Models																			
Lingshu	7B	58.94	78.55	39.89	54.29	58.24	8.56	45.28	58.73	51.52	50.44	52.60	49.84	53.41	48.95	53.76	49.60	43.69	50.26
HuatuoGPT-V	7B	45.74	71.39	43.09	39.05	35.11	9.61	47.73	45.24	49.09	42.89	51.86	51.47	55.68	49.20	55.08	52.93	45.78	51.71
Qwen3-VL (SFT)	**8B**	72.34	76.73	65.74	78.10	71.01	59.76	75.87	53.97	60.41	68.21	94.43	**95.38**	94.57	91.70	95.28	88.36	57.96	88.24
BreastGPT (cluster)	**8B**	**86.81**	77.21	**75.00**	**82.86**	77.13	**68.32**	**81.12**	**61.11**	**71.38**	**75.66**	**95.97**	95.29	**95.48**	93.24	**95.72**	**89.93**	**63.80**	**89.92**
BreastGPT (learn)	**8B**	84.47	71.03	68.51	75.71	**78.46**	55.26	73.95	57.14	71.25	70.64	95.86	94.73	87.12	**93.57**	85.22	81.78	63.36	85.95

Table 1: VQA performance on BreastStage-Bench. The left half reports **closed-ended accuracy (%)** and the right half reports **open-ended normalized score (%)** for the same set of models. Open-ended evaluation does not include Mammography. Avg is the mean across the task columns of the corresponding half.

in the extended pool we add in the supplement (MedDr (He et al., 2024), Hulu-Med (Jiang et al., 2025), LLaVA-OneVision-1.5 (An et al., 2025)) trail by more than 25 points. Generation results tell the same story with somewhat tighter margins, because the weighted metric saturates near the top. The one place BreastGPT stands out qualitatively rather than quantitatively is 3D grounding on CT and MRI, where it is the only model in the pool that emits non-trivial volumetric boxes at all. Per-modality IoU and the eight extended baselines (Grok-4.1-Fast (xAI, 2025), Gemma-4 (Team et al., 2024), GLM-4.6V-Flash (Team et al., 2025b), LLaVA-OneVision-1.5, HealthGPT (Lin et al., 2025b), Hulu-Med, MedDr, RadFM (Wu et al., 2025)) are in § Tables 10, 11 and 13. **Where the baselines break.** Frontier proprietary models are not particularly close: the Claude and Gemini families sit in the low forties to low fifties on closed-ended VQA, and even GPT-5.4 trails BreastGPT by more than twenty points. Mammography-based diagnosis is where they fall apart the hardest, with the Claude family scoring in the single digits to high teens. We read this as evidence that BI-RADS reading, molecular subtyping, and multiparametric MRI interpretation are not the kind of skill that gets picked up from generic web pretraining. The medical-specific baselines do not rescue this either: Lingshu and HuatuoGPT-V land near or below the general-purpose 7–8B VLMs, so medical pretraining corpora as they stand do not transfer to breast workflow tasks.

		Caption							Report
		BUS		CT	Mammo		Histo	MRI	MRI
Model	#P	IoU	Wtd	Wtd	IoU	Wtd	Wtd	Wtd	Wtd
Proprietary Models									
GPT-5.4	–	6.68	47.66	43.95	9.08	45.10	44.36	47.00	55.16
Claude-opus-4-6	–	11.41	46.99	43.68	12.06	47.04	44.94	47.75	48.56
Claude-sonnet-4-6	–	7.40	46.60	42.82	–	44.79	44.21	49.30	49.30
Gemini-3.1-Flash	–	56.20	51.48	46.01	7.52	49.29	44.89	51.45	51.03
Gemini-3.1-Pro	–	–	45.60	44.64	–	42.36	45.84	54.86	54.86
Open-Source Models									
Qwen2.5-VL	3B	2.67	46.84	46.88	1.38	48.36	45.71	50.69	49.59
Qwen2.5-VL	7B	6.38	46.79	47.31	1.47	48.11	45.07	51.95	50.71
Qwen3-VL	4B	33.64	46.68	41.66	4.29	46.19	45.38	37.66	49.46
Qwen3-VL	8B	29.35	46.97	47.91	4.06	47.09	46.03	49.57	50.49
MiMo-VL	7B	12.35	43.22	46.14	2.22	43.80	43.98	47.32	46.55
InternVL3.5	8B	0.35	46.97	47.69	0.17	45.72	45.77	49.28	49.12
Medical-Specific Models									
Lingshu	7B	5.00	48.43	49.45	0.55	49.14	46.86	51.84	51.14
HuatuoGPT-V	7B	0.21	47.46	48.92	0.47	48.84	46.24	51.46	51.00
Qwen3-VL (SFT)	**8B**	72.79	76.61	72.72	8.62	76.68	62.42	65.21	66.10
BreastGPT (cluster)	**8B**	**79.59**	**79.32**	73.16	**23.14**	**77.64**	66.78	**67.69**	**67.67**
BreastGPT (learn)	**8B**	–	79.16	**77.07**	–	76.54	**68.11**	–	66.04

Table 2: Caption and Report Generation Performance (%). IoU = mean grounding IoU on the grounded captioning task, "-" indicates the model produces no usable bbox in that modality. Wtd = weighted composite ($0.5\times$BERT $+ 0.25\times$BLEU $+ 0.25\times$R-1). IoU on the 3D modalities (CT, MRI) is reported in Table 13.

Cluster vs. learn. The two BreastGPT variants differ only in how the final 128 tokens are picked. The training-free cluster selector comes out ahead by five points on closed-ended VQA and four on open-ended, without adding a single parameter. We read this as the coverage objective already encoding most of what a learned selector would need to discover, so the cross-attention selector ends up paying an optimization cost without much to recoup at this scale.

5.4 Ablation Studies

Does the GigaPixel branch matter? We retrain BreastGPT without the GigaPixel branch, so WSIs

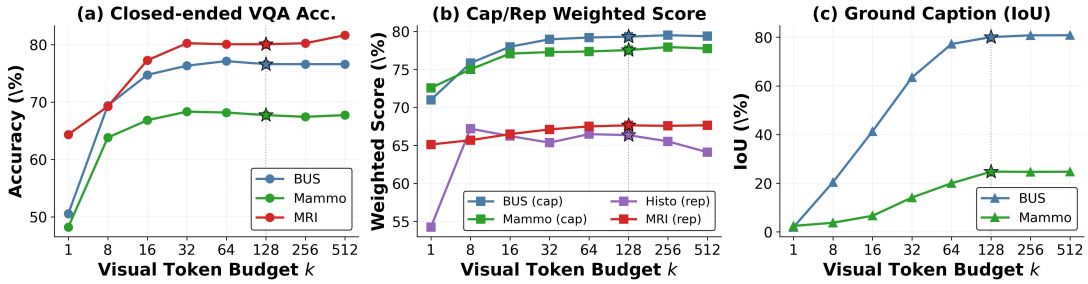


Figure 4: Visual token budget sweep on BreastStage-Bench. **(a)** Closed-ended VQA accuracy on diagnosis-stage tasks for BUS, mammography, and MRI. **(b)** Caption and report weighted score for BUS, mammography, histopathology report, and MRI report. **(c)** Ground caption IoU on BUS and mammography. Dotted vertical lines and stars mark the chosen operating point $k = 128$.

go through the standard ViT branch with 32 random 512×512 tiles, holding everything else (backbone, data, recipe) constant (Qwen3-VL (SFT) row in Sections 5 and 5.3). On radiology the gap to BreastGPT is small, which is expected given that the source images already fit a ViT. Histopathology is where the branch becomes important: a 32-tile random window covers less than one percent of a typical breast WSI, so tissue architecture and rare regions tied to subtype, grade, and response are largely missed, and the closed-ended gap to BreastGPT grows to roughly eleven points. The CONCH+LongNet pipeline plus coverage selection is what substantially reduces this gap. Per-component breakdown is in § Table 16.

How many visual tokens do we need? We sweep $k \in \{1, 8, 16, 32, 64, 128, 256, 512\}$ across five modalities and three task families (Fig. 4; full numbers in § Tables 17 and 18). Closed-ended VQA flattens out early: radiology plateaus by k=32, and BUS-Diagnosis is essentially at its final accuracy from k=64 on. Generation peaks near k=128 and from there yields almost nothing for radiology; on histopathology, where most of the slide is redundant tissue, the score actually drifts down as we keep adding tokens. Grounding is the only task that genuinely wants a larger budget, since a single token cannot even represent a box, but even there k=128 recovers over 99% of what k=512 gets. We therefore use k=128 throughout the paper: it sits on the plateau for VQA and generation, keeps grounding nearly optimal, and lets radiology and pathology share the same downstream cost.

GPU time. On a representative WSI containing approximately six thousand patches, inference with k=128 takes around 200 ms, with the selector contributing less than 10 ms. Feeding all patch tokens directly into the LLM (i.e. without token selection) increases latency by more than an order of magnitude (approximately 33×). Increasing the token budget to k=512 further adds several hundred milliseconds while yielding less than one point of performance gain. These results indicate that the selector introduces only a small constant overhead while substantially reducing the dominant inference cost. Detailed per-k latency measurements are provided in § Table 15 and § Fig. 5.

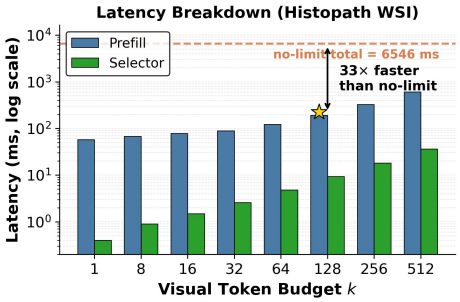


Figure 5: GPU latency on a histopathology WSI. Two bars per k give prefill and selector latency. The dashed line is the no-limit baseline that bypasses the selector. Log-scaled y-axis.

6 Conclusion

Breast cancer care requires reasoning across screening, diagnosis, and treatment, yet existing medical MLLMs remain largely optimized for isolated stages. To bridge this gap, we introduce **BreastStage**, a workflow-aligned multimodal dataset and benchmark, and **BreastGPT**, a unified VLM that combines stage-aware role prompting with a dual-branch encoder for cross-scale modeling from radiology to gigapixel pathology. Experiments on BreastStage-Bench show that our workflow-aligned design substantially narrows the cross-stage reasoning gap where both general-purpose and medical-specific MLLMs struggle. We hope this work encourages future medical multimodal models to move beyond isolated tasks toward workflow-centred reasoning.

Limitations

While BreastStage covers screening, diagnosis, and treatment planning, most training samples are collected from different patient cohorts rather than from longitudinal records of the same patients. Therefore, the current training corpus supports workflow-aligned supervision at the task and modality level, but does not fully capture patient-level temporal continuity across the entire care pathway. Notably, a subset in BreastStage-bench includes full-workflow cases from the same patients, allowing us to partially examine whether the model can integrate clinically connected evidence across stages. Building a per-patient cross-stage corpus is a field-level bottleneck, requiring linked records across departments, multi-year follow-up windows, and IRB approval for identifiable longitudinal data and we are working for this.

Acknowledgments

This work was supported in part by the Postdoctoral Research Fund of West China Hospital, Sichuan University (Project No. 2025HXBH040); the Sichuan Science and Technology Program (Grant No. 2026NSFSC1775); the National Natural Science Foundation of China (Grant No. 62476246); the "Pioneer" and "Leading Goose" R&D Program of Zhejiang Province (Grant No. 2025C02120); and Guangzhou City's Key R&D Program of China (Grant No. 2024B01J1301). Additional support was provided by DAMO Academy through the DAMO Academy Research Intern Program. Jiajin Zhang was also supported in part by the Zhejiang Province Postdoctoral Research Excellence Funding Program (Grant No. ZJ2025048).

References

Faruk Ahmed, Andrew Sellergren, Lin Yang, Shawn Xu, Boris Babenko, Abbi Ward, Niels Olson, Arash Mohtashamian, Yossi Matias, Greg S Corrado, and 1 others. 2024. Pathalign: A vision-language model for whole slide images in histopathology. *arXiv preprint arXiv:2406.19578*.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, and 1 others. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.

Rawan AlSaad, Alaa Abd-Alrazaq, Sabri Boughorbel, Arfan Ahmed, Max-Antoine Renault, Rafat Damseh, and Javaid Sheikh. 2024. Multimodal large language models in health care: applications, challenges, and future outlook. *Journal of medical Internet research*, 26:e59505.

Asmaa S Alsolami, Wafaa Shalash, Wafaa Alsaggaf, Sawsan Ashoor, Haneen Refaat, and Mohammed Elmogy. 2021. King abdulaziz university breast cancer mammogram dataset (kau-bcmd). *Data*, 6(11):111.

Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Didi Zhu, and 1 others. 2025. Llava-onevision-1.5: Fully open framework for democratized multimodal training. *arXiv preprint arXiv:2509.23661*.

N Aristokli, I Polycarpou, SC Themistocleous, D Sophocleous, and I Mamais. 2022. Comparison of the diagnostic performance of magnetic resonance imaging (mri), ultrasound and mammography for detection of breast cancer based on tumor type, breast density and patient's history: A review. *Radiography*, 28(3):848–856.

Melina Arnold, Eileen Morgan, Harriet Rumgay, Allini Mafra, Deependra Singh, Mathieu Laversanne, Jerome Vignat, Julie R Gralow, Fatima Cardoso, Sabine Siesling, and 1 others. 2022. Current and future burden of breast cancer: Global statistics for 2020 and 2040. *The Breast*, 66:15–23.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, and 1 others. 2025a. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, and 1 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Zak Buzzard, Konstantin Hemker, Nikola Simidjievski, and Mateja Jamnik. 2024. Paths: A hierarchical transformer for efficient whole slide image analysis. *arXiv preprint arXiv:2411.18225*.

Chris Carr, PhD FelipeKitamura, MD, George Partridge, inversion, Jayashree Kalpathy-Cramer, John Mongan, Katherine Andriole, Lavender, Maryam Vazirabad, Michelle Riopel, Robyn Ball, Sohier Dane, and Yan Chen. 2022. Rsna screening mammography breast cancer detection. `https://kaggle.com/competitions/rsna-breast-cancer-detection`. Kaggle.

Jintai Chen, Yaojun Hu, Mingchen Cai, Yingzhou Lu, Yue Wang, Xu Cao, Miao Lin, Hongxia Xu, Jian Wu, Xiao Cao, and 1 others. 2025. Trialbench: Multimodal ai-ready datasets for clinical trial prediction. *Scientific data*, 12(1):1564.

Junying Chen, Chi Gui, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong Wang, Ruifei Zhang, Zhenyang Cai, Ke Ji, and 1 others. 2024a. Huatuogpt-vision, towards injecting medical visual knowledge into multimodal llms at scale. *arXiv preprint arXiv:2406.19280*.

Pengcheng Chen, Jin Ye, Guoan Wang, Yanjun Li, Zhongying Deng, Wei Li, Tianbin Li, Haodong Duan, Ziyan Huang, Yanzhou Su, and 1 others. 2024b. Gmai-mmbench: A comprehensive multimodal evaluation benchmark towards general medical ai. *Advances in Neural Information Processing Systems*, 37:94327–94427.

Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Bowen Chen, Andrew Zhang, Daniel Shao, Andrew H Song, Muhammad Shaban, and 1 others. 2023. A general-purpose self-supervised model for computational pathology. *arXiv preprint arXiv:2308.15474*.

Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Andrew H Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, and 1 others. 2024c. Towards a general-purpose foundation model for computational pathology. *Nature Medicine*, 30(3):850–862.

Jiayu Ding, Shuming Ma, Li Dong, Xingxing Zhang, Shaohan Huang, Wenhui Wang, Nanning Zheng, and Furu Wei. 2023. Longnet: Scaling transformers to 1,000,000,000 tokens. *arXiv preprint arXiv:2307.02486*.

Sixun Dong, Juhua Hu, Mian Zhang, Ming Yin, Yanjie Fu, and Qi Qian. 2026. MMtok: Multimodal coverage maximization for efficient inference of VLMs. In *The International Conference on Learning Representations*.

Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, and 1 others. 2024. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the ACM International Conference on Multimedia*, pages 11198–11201.

Omar Freihat, David Sipos, and Arpad Kovacs. 2025. Global burden and projections of breast cancer incidence and mortality to 2050: a comprehensive analysis of globocan data. *Frontiers in Public Health*, 13:1622954.

Zhengrui Guo, Qichen Sun, Jiabo Ma, Lishuang Feng, Jinzhuo Wang, and Hao Chen. 2025. Context matters: Query-aware dynamic long sequence modeling of gigapixel images. *arXiv preprint arXiv:2501.18984*.

Ibrahim Ethem Hamamci, Sezgin Er, Chenyu Wang, Furkan Almas, Ayse Gulnihan Simsek, Sevval Nil Esirgun, Irem Dogan, Omer Faruk Durugol, Benjamin Hou, Suprosanna Shit, and 1 others. 2026. Generalist foundation models from a multimodal dataset for 3d computed tomography. *Nature Biomedical Engineering*, pages 1–19.

Jing Hao, Yuxuan Fan, Yanpeng Sun, Kaixin Guo, Lizhuo Lin, Jinrong Yang, Qi Yong H Ai, Lun M Wong, Hao Tang, and Kuo Feng Hung. 2025. Towards better dental ai: A multimodal benchmark and instruction dataset for panoramic x-ray analysis. *Advances in Neural Information Processing Systems*.

Jing Hao, Yuci Liang, Lizhuo Lin, Yuxuan Fan, Wenkai Zhou, Kaixin Guo, Zanting Ye, Yanpeng Sun, Xinyu Zhang, Yanqi Yang, and 1 others. 2026. Oralgpt-omni: A versatile dental multimodal large language model. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Iryna Hartsock and Ghulam Rasool. 2024. Vision-language models for medical report generation and visual question answering: A review. *Frontiers in artificial intelligence*, 7:1430984.

Sunan He, Yuxiang Nie, Zhixuan Chen, Zhiyuan Cai, Hongmei Wang, Shu Yang, and Hao Chen. 2024. Meddr: Diagnosis-guided bootstrapping for large-scale medical vision-language learning. *arXiv e-prints*, pages arXiv–2404.

Yaojun Hu, Danyang Tu, Yang Liu, Jiajin Zhang, Wei Fang, Zhiqiang Liu, Chunlai Dong, Yingda Xia, Haochao Ying, Jian Wu, and Ling Zhang. 2026. Sever: Selective visual exposure and retrieval for 3d medical image question answering. *Preprint*, arXiv:2608.25630.

Yutao Hu, Tianbin Li, Quanfeng Lu, Wenqi Shao, Junjun He, Yu Qiao, and Ping Luo. 2024. Omnimedvqa: A new large-scale comprehensive evaluation benchmark for medical lvlm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22170–22183.

Fuxiang Huang, Jiayi Zhu, Yunfang Yu, Yu Xie, Yuan Guo, Qingcong Kong, Mingxiang Wu, Xinrui Jiang, Shu Yang, Jiabo Ma, and 1 others. 2025. A versatile foundation model for ai-enabled mammogram interpretation. *arXiv preprint arXiv:2509.20271*.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.

Jiwoong J Jeong, Brianna L Vey, Ananth Bhimireddy, Thomas Kim, Thiago Santos, Ramon Correa, Raman Dutt, Marina Mosunjac, Gabriela Oprea-Ilies, Geoffrey Smith, and 1 others. 2023. The emory breast imaging dataset (embed): A racially diverse, granular dataset of 3.4 million screening and diagnostic mammographic images. *Radiology: Artificial Intelligence*, 5(1):e220047.

Songtao Jiang, Yuan Wang, Sibo Song, Tianxiang Hu, Chenyi Zhou, Bin Pu, Yan Zhang, Zhibo Yang, Yang Feng, Joey Tianyi Zhou, and 1 others. 2025. Hulumed: A transparent generalist model towards holistic medical vision-language understanding. *arXiv preprint arXiv:2510.08668*.

Béria Chingnabé Kalpélbé, Angel Gabriel Adaambiik, and Wei Peng. 2025. Vision language models in medicine. *arXiv preprint arXiv:2503.01863*.

Laskarina Katsika, Eirini Boureka, Ioannis Kalogiannidis, Ioannis Tsakiridis, Ilias Tirodimos, Konstantinos Lallas, Zoi Tsimtsiou, and Themistoklis Dagklis. 2024. Screening for breast cancer: a comparative review of guidelines. *Life*, 14(6):777.

Rana Khaled, Maha Helal, Omar Alfarghaly, Omnia Mokhtar, Abeer Elkorany, Hebatalla El Kassas, and Aly Fahmy. 2022. Categorized contrast enhanced mammography dataset for diagnostic and artificial intelligence research. *Scientific data*, 9(1):122.

Joanne Kim, Andrew Harper, Valerie McCormack, Hyuna Sung, Nehmat Houssami, Eileen Morgan, Miriam Mutebi, Gail Garvey, Isabelle Soerjomataram, and Miranda M Fidler-Benaoudia. 2025. Global patterns and trends in breast cancer incidence and mortality across 185 countries. *Nature Medicine*, 31(4):1154–1162.

Rebecca Sawyer Lee, Francisco Gimenez, Assaf Hoogi, Kanae Kawai Miyake, Mia Gorovoy, and Daniel L Rubin. 2017. A curated mammography data set for use in computer-aided detection and diagnosis research. *Scientific data*, 4(1):170177.

Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2023. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36:28541–28564.

Honglin Li, Yunlong Zhang, Pingyi Chen, Zhongyi Shui, Chenglu Zhu, and Lin Yang. 2024. Rethinking transformer for long contextual histopathology whole slide image analysis. *Advances in Neural Information Processing Systems*, 37:101498–101528.

Haoneng Lin, Cheng Xu, and Jing Qin. 2025a. Taming vision-language models for medical image analysis: A comprehensive review. *arXiv preprint arXiv:2506.18378*.

Tianwei Lin, Wenqiao Zhang, Sijing Li, Yuqian Yuan, Binhe Yu, Haoyuan Li, Wanggui He, Hao Jiang, Mengze Li, Xiaohui Song, and 1 others. 2025b. Healthgpt: A medical large vision-language model for unifying comprehension and generation via heterogeneous knowledge adaptation. *arXiv preprint arXiv:2502.09838*.

Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, and 1 others. 2024. Mmbench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pages 216–233. Springer.

Kosmia Loizidou, Galateia Skouroumouni, Costas Pitris, and Christos Nikolaou. 2021. Digital subtraction of temporally sequential mammograms for improved detection and classification of microcalcifications. *European radiology experimental*, 5(1):40.

Ming Y Lu, Bowen Chen, Drew FK Williamson, Richard J Chen, Ivy Liang, Tong Ding, Guillaume Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, and 1 others. 2024. A visual-language foundation model for computational pathology. *Nature Medicine*, 30(3):863–874.

Luyang Luo, Mingxiang Wu, Mei Li, Yi Xin, Qiong Wang, Varut Vardhanabhuti, Winnie CW Chu, Zhenhui Li, Juan Zhou, Pranav Rajpurkar, and 1 others. 2025. A large model for non-invasive and personalized management of breast cancer from multiparametric mri. *Nature communications*, 16(1):3647.

Christine McCaffrey, Chowdhury Jahangir, Clodagh Murphy, Caoimbhe Burke, William M Gallagher, and Arman Rahman. 2024. Artificial intelligence in digital histopathology for predicting patient prognosis and treatment efficacy in breast cancer. *Expert review of molecular diagnostics*, 24(5):363–377.

Yu Meng, Kaiyuan Li, Chenran Huang, Chen Gao, Xinlei Chen, Yong Li, and Xiaoping Zhang. 2025. Plphp: Per-layer per-head vision token pruning for efficient large vision-language models. *arXiv preprint arXiv:2502.14504*.

Inês C Moreira, Igor Amaral, Inês Domingues, António Cardoso, Maria Joao Cardoso, and Jaime S Cardoso. 2012. Inbreast: toward a full-field digital mammographic database. *Academic radiology*, 19(2):236–248.

Hieu T Nguyen, Ha Q Nguyen, Hieu H Pham, Khanh Lam, Linh T Le, Minh Dao, and Van Vu. 2023. Vindr-mammo: A large-scale benchmark dataset for computer-aided diagnosis in full-field digital mammography. *Scientific Data*, 10(1):277.

Yasmine Omri, Parth Shroff, and Thierry Tambe. 2025. Token sequence compression for efficient multimodal computing. *arXiv preprint arXiv:2504.17892*.

Parita Oza, Urvi Oza, Rajiv Oza, Paawan Sharma, Samir Patel, Pankaj Kumar, and Bakul Gohel. 2024. Digital mammography dataset for breast cancer diagnosis research (dmid) with breast mass segmentation analysis. *Biomedical Engineering Letters*, 14(2):317–330.

Jiao Qu, Yang Liu, Mengmei Zhang, Huan Xie, Juan Yang, Mengyuan Xu, Ling Yang, Wenlei Qian, Yan Li, Jing Jiang, and 1 others. 2026. Multiparameter concept-based interpretable model for early breast cancer diagnosis and structured reporting: a multicenter, multi-reader, radiologist-in-the-loop study. *BMC medicine*.

Rui Sha, Xiang-meng Kong, Xin-yu Li, and Ya-bing Wang. 2024. Global burden of breast cancer and attributable risk factors in 204 countries and territories, from 1990 to 2021: results from the global burden of disease study 2021. *Biomarker Research*, 12(1):87.

Yuzhang Shang, Mu Cai, Bingxin Xu, Yong Jae Lee, and Yan Yan. 2025. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22857–22867.

Moein Sorkhei, Yue Liu, Hossein Azizpour, Edward Azavedo, Karin Dembrower, Dimitra Ntoula, Athanasios Zouzos, Fredrik Strand, and Kevin Smith. 2021. Csaw-m: An ordinal classification dataset for benchmarking mammographic masking of cancer. *arXiv preprint arXiv:2112.01330*.

John Suckling. 1994. The mammographic images analysis society digital mammogram database. In *Exerpta Medica International Congress*, volume 1069, pages 375–378.

Yizheng Sun, Yanze Xin, Hao Li, Jingyuan Sun, Chenghua Lin, and Riza Theresa Batista-Navarro. 2025. Lvpruning: An effective yet simple language-guided vision token pruning approach for multi-modal large language models. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4299–4308.

Hao Tang and Chengchao Shen. 2025. Learning compact vision tokens for efficient large multimodal models. *arXiv*.

Core Team, Zihao Yue, Zhenru Lin, Yifan Song, Weikun Wang, and 1 others. 2025a. Mimo-vl technical report. *Preprint*, arXiv:2506.03569.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, and 1 others. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.

V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, and 1 others. 2025b. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning. *Preprint*, arXiv:2507.01006.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Hui Hui, Yanfeng Wang, and Weidi Xie. 2025. Towards generalist foundation model for radiology by leveraging web-scale 2d&3d medical data. *Nature Communications*, 16(1):7866.

xAI. 2025. Grok 4 fast model card. https://data.x.ai/2025-09-19-grok-4-fast-model-card.pdf.

Hanguang Xiao, Feizhong Zhou, Xingyue Liu, Tianqi Liu, Zhipeng Li, Xin Liu, and Xiaoxuan Huang. 2025. A comprehensive survey of large language models and multimodal large language models in medicine. *Information Fusion*, 117:102888.

Conghao Xiong, Hao Chen, and Joseph JY Sung. 2025. A survey of pathology foundation model: Progress and future directions. *arXiv preprint arXiv:2504.04045*.

Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, and 1 others. 2024. A whole-slide foundation model for digital pathology from real-world data. *Nature*, 630(8015):181–188.

Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Ruifeng Yuan, Jianyu Wang, Chenghao Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, and 1 others. 2025. Lingshu: A generalist foundation model for unified multimodal medical understanding and reasoning. *arXiv preprint arXiv:2506.07044*.

Jiarui Ye and Hao Tang. 2025. Multimodal large language models for medicine: A comprehensive survey. *arXiv preprint arXiv:2504.21051*.

Haojun Yu, Youcheng Li, Zihan Niu, Nan Zhang, Xuantong Gong, Huan Li, Zhiying Zou, Haifeng Qi, Zhenxiao Cao, Zijie Lan, Xingjian Yuan, Jiating He, Haokai Zhang, Shengtao Zhang, Zicheng Wang, Dong Wang, Ziwei Zhao, Congying Chen, Yong Wang, and 2 others. 2026. A chain-of-thought reasoning breast ultrasound dataset covering all histopathology categories. *Scientific Data*, 13:236.

Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.

Weihao Yu, Zhengyuan Yang, Lingfeng Ren, Linjie Li, Jianfeng Wang, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, and Xinchao Wang. 2024. Mm-vet v2: A challenging benchmark to evaluate large multimodal models for integrated capabilities. *arXiv preprint arXiv:2408.00765*.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi.

In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.

Jiaheng Zhou, Wei Fang, Luyuan Xie, Yanfeng Zhou, Lianyan Xu, Minfeng Xu, Ge Yang, and Yuxing Tang. 2026. Dual-res tandem mamba-3d: Bilateral breast lesion detection and classification on non-contrast chest CT. In *Advances in Neural Information Processing Systems*.

Jiayi Zhu, Fuxiang Huang, Qiong Luo, and Hao Chen. 2025. A benchmark for breast cancer screening and diagnosis in mammogram visual question answering. *Nature Communications*, 16(1):11683.

Contents

A Appendix Overview **13**
A.1 LLM Usage Statement 13
A.2 Ethics Approval and Privacy . . . 13
A.3 Additional Limitations 14
A.4 Compute Resources 14
A.5 Code and Data Availability 15

B Related Work **15**

C BreastStage Curation Details **16**
C.1 Data Sources and Modality Coverage 16
C.2 Instruction Data Generation Pipeline 17
C.3 Task Taxonomy and Distribution . 20
C.4 Human Annotator Statement . . . 21

D BreastGPT Technical Details **21**
D.1 Base Model Configuration 21
D.2 GigaPixel Encoder Implementation 22
D.3 Coverage-Maximizing Token Selection 24
D.4 Training Configuration 27
D.5 System Prompt Templates 27

E BreastStage-Bench Evaluation **27**
E.1 Benchmark Construction 27
E.2 Evaluation Metrics 27
E.3 Breast Expert Validation 29
E.4 Baseline Model Configuration . . 30

F Additional Results and Ablations **30**
F.1 Per-Modality Performance Breakdown 30
F.2 Extended Baseline Comparisons . 30
F.3 Open-ended VQA Raw Metric Breakdown 30
F.4 3D Grounding IoU on CT and MRI 30
F.5 BreastGPT Grounding Recognition Breakdown 30
F.6 Inference Efficiency Analysis . . . 30

G Detailed Ablation Results **32**
G.1 Numerical Results for WSI Branch Ablation 32
G.2 Numerical Results for Visual Token Budget Sweep 32

H Qualitative Analysis and Case Studies **33**
H.1 Case Analysis 33

A Appendix Overview

This Appendix material provides the technical and procedural details that support the main claims of the paper. We first describe the construction of **BreastStage**, including data provenance, modality coverage, instruction generation, quality control, and task taxonomy. We then provide implementation details for **BreastGPT**, with particular emphasis on the GigaPixel WSI branch and the coverage-maximizing visual token selector that enables a unified token budget across standard radiology images and gigapixel pathology slides. Finally, we describe the BreastStage-Bench evaluation protocol, additional quantitative analyses, qualitative examples, limitations, ethics, and release plans.

This Appendix is intended to make three aspects of the paper auditable: (i) how heterogeneous breast imaging data are mapped into a stage-aware clinical workflow, (ii) how BreastGPT processes image scales that differ by several orders of magnitude, and (iii) how benchmark results, ablations, and error analyses are produced.

A.1 LLM Usage Statement

In this work, we primarily employ LLMs in two aspects: (i) data construction, where Qwen2.5-VL-72B and Qwen3-Max drive the orchestration pipeline described in Appendix C.2; and (ii) manuscript polishing, where we use LLMs to improve grammar and clarity of the writing without altering the technical claims.

A.2 Ethics Approval and Privacy

BreastStage combines two kinds of sources with different ethical pathways. The BUS, mammography, CT, and histopathology subsets come entirely from public datasets that were de-identified by their original publishers and released under their respective data-use terms; we use them within those terms

and do not perform any additional re-identification of individual samples. The MRI subset is the only institutional cohort: it was acquired from two collaborating hospitals [anonymous for review] under IRB approval, with patient consent obtained where required by the contributing institutions.

De-identification of the institutional MRI cohort. Because the MRI cohort is the only data source that originates inside hospitals, our de-identification protocol is scoped specifically to it. Before MRI data leave either source institution, identifying information is removed at three levels:

- **DICOM header level.** Patient names, dates of birth, medical-record numbers, hospital identifiers, referring-physician fields, and exam dates are stripped; absolute timestamps are converted to relative offsets so longitudinal ordering is preserved without disclosing calendar dates.
- **Image content level.** The MRI volumes are reconstructed within a breast-coil field of view that does not include the patient's face, so no additional image-side scrubbing is required.
- **Free-text level.** The accompanying Chinese radiology and pathology reports are passed through the bilingual parsers in Figure 8, which emit a strictly schema-bounded English JSON so that no free-text PHI (referring physician name, hospital department, etc.) propagates into the released instruction-following pairs.

Access control. Access to identifiable raw MRI data is restricted to authorised research personnel under signed data-use agreements. The public mirror of BreastStage will release only the de-identified, instruction-formatted records and the derived bounding-box / mask annotations; raw DICOMs from the institutional MRI cohort are not redistributed. The other four subsets (BUS, mammography, CT, histopathology) are accessed via their original public-dataset distribution channels.

Risks acknowledged. Because the model is trained on clinical imaging data, performance may vary across institutions, scanners, demographic groups, and disease subtypes; we recommend that any downstream use perform site-specific validation before clinical deployment. The benchmark is intended to encourage transparent comparison of breast cancer workflow models, not to replace clinician judgement.

A.3 Additional Limitations

Dataset. While BreastStage covers five major imaging modalities routinely used in breast cancer screening and diagnosis, it does not currently include molecular imaging modalities such as PET/CT, which are often used for staging, metastasis assessment, and treatment response evaluation in advanced disease. Emerging modalities including photoacoustic imaging are also beyond the current scope. In addition, the present workflow primarily focuses on imaging-centered decision support, while longitudinal clinical information such as treatment trajectories, genomic assays, laboratory tests, and medication records remains only partially available across cohorts. Future work should extend BreastStage toward longitudinal multimodal patient records and investigate whether workflow-aligned MLLMs can support temporal clinical reasoning for treatment response evaluation, recurrence monitoring, and survivorship management.

Clinical deployment. BreastGPT is currently developed as a research prototype and has not been clinically validated or reviewed by regulatory authorities. Any use in real-world clinical workflows would require prospective clinical evaluation, compliance with applicable regulatory and data-governance requirements, integration with existing hospital information systems, and continuous post-deployment surveillance to monitor model drift, performance degradation, and potential safety risks. Accordingly, BreastGPT should be regarded as a research and benchmarking foundation model, not as an autonomous diagnostic or treatment-planning system.

A.4 Compute Resources

Model training. BreastGPT training was carried out on 32 NVIDIA H100 GPUs (4 nodes $\times$ 8 GPUs) with DeepSpeed ZeRO-2, FlashAttention, and bfloat16 precision. End-to-end wall-clock training time across the two stages was 3 days, 7 hours, 19 minutes, 53 seconds, corresponding to approximately 2,530 H100-hours.

Dataset construction and evaluation. Visual quality control with Qwen2.5-VL-72B and Qwen3-Max-driven instruction generation are run through external APIs and shared infrastructure; per-call

API expenditure was not separately tracked. Specialist annotation of MRI tumours (10 board-certified breast radiologists) and the 3-specialist Bench audit were conducted by clinical collaborators rather than billed compute.

A.5 Code and Data Availability

Code. An anonymised snapshot of the BreastGPT codebase, including training scripts, evaluation pipelines, and the BreastStage data-construction tooling, is available for review at `https://anonymous.4open.science/r/BreastGPT`. The full code release with model checkpoints will be moved to a non-anonymised public repository upon publication. Pre-trained BreastGPT model weights will be released on Hugging Face.

Data. The BreastStage dataset and BreastStage-Bench evaluation suite will be made publicly available. For components derived from institutional clinical data, release will follow the corresponding IRB protocol, de-identification requirements, and data-use agreements. When raw image release is restricted, we will provide metadata, task definitions, evaluation code, and access instructions. The BreastStage dataset, BreastStage-Bench, and BreastGPT weights will be released under the **CC BY-NC 4.0** license; the BreastGPT code under **Apache 2.0**.

B Related Work

Recent medical MLLMs have extended general-purpose vision-language modeling into clinical image understanding, report generation, visual question answering, and interactive decision support (Ye and Tang, 2025; Kalpélbé et al., 2025; AlSaad et al., 2024; Xiao et al., 2025; Lin et al., 2025a; Hartsock and Rasool, 2024; Qu et al., 2026; Chen et al., 2025). Representative systems such as LLaVA-Med, HuatuoGPT-Vision, MedDr, HealthGPT, and Lingshu show that biomedical instruction tuning can improve clinical-domain perception and dialogue (Li et al., 2023; Chen et al., 2024a; He et al., 2024; Lin et al., 2025b; Xu et al., 2025). Meanwhile, broad evaluation suites such as OmniMedVQA, GMAI-MMBench, MMMU, MMBench, and VLMEvalKit expose the gap between generic multimodal ability and reliable medical reasoning (Hu et al., 2024; Chen et al., 2024b; Yue et al., 2024; Liu et al., 2024; Duan et al., 2024). Domain-specific studies in dentistry further show that clinically meaningful evaluation often requires benchmarks organized around specialty-specific anatomy, modalities, reporting conventions, and treatment workflows (Hao et al., 2025, 2026). BreastGPT follows this specialty-grounded direction for breast oncology, where the challenge is not only domain-specific perception but also workflow-aware reasoning across screening, diagnosis, and treatment.

Breast imaging AI has progressed rapidly, yet most prior work remains modality-specific. In mammography, VersaMammo builds a large multi-institutional mammogram foundation model and evaluates detection, segmentation, classification, retrieval, and VQA tasks (Huang et al., 2025), while MammoVQA standardizes mammogram VQA across 15 public datasets and shows that both general-purpose and medical-specific LVLMs remain unreliable for mammogram interpretation (Zhu et al., 2025). For breast ultrasound, BUS-CoT provides chain-of-thought reasoning annotations across all 99 WHO histopathology categories, connecting observations, imaging features, BI-RADS assessment, pathology labels, and reasoning traces (Yu et al., 2026). For MRI, MOME integrates multiparametric breast MRI through a mixture-of-modality-experts design and supports malignancy identification, biopsy recommendation, triple-negative breast cancer classification, and neoadjuvant chemotherapy response prediction (Luo et al., 2025). These works show the value of breast-specific modeling, but they do not by themselves provide a single model or benchmark spanning the full clinical pathway.

Computational pathology has similarly moved from patch-level representation learning toward slide-level and vision-language foundation models. Recent pathology foundation models, including UNI, CONCH, and large-scale whole-slide encoders, learn transferable morphology representations for diagnostic, prognostic, retrieval, and report-generation tasks (Chen et al., 2024c; Lu et al., 2024; Chen et al., 2023; Xu et al., 2024). WSI vision-language and long-sequence modeling efforts such as PathAlign, query-aware gigapixel modeling, hierarchical WSI transformers, and long-context histopathology transformers further bridge patch-level embeddings with slide-level reasoning and language interaction (Ahmed et al., 2024; Guo et al., 2025; Buzzard et al., 2024; Li et al., 2024; Xiong et al., 2025). These studies highlight a central difficulty for BreastGPT: pathology slides con-

tain sparse, distributed, and scale-dependent evidence, while breast radiology images use much smaller visual inputs and different clinical semantics.

Efficient visual token selection is therefore essential for unified breast workflow modeling. Recent methods reduce MLLM inference cost through token pruning, learned token selection, adaptive compression, query-aware retrieval, and layer-wise token reduction (Dong et al., 2026; Shang et al., 2025; Tang and Shen, 2025; Sun et al., 2025; Meng et al., 2025; Omri et al., 2025; Hu et al., 2026). MMTok formulates token selection as multimodal coverage maximization, balancing text-vision alignment with vision-vision diversity (Dong et al., 2026). BreastGPT adapts this principle to clinical workflow modeling: for radiology images, the selector removes redundant visual tokens while preserving prompt-relevant lesion and anatomical information; for WSIs, it compresses LongNet-contextualized patch embeddings into a fixed-size representation that retains both query-relevant pathology regions and global tissue coverage.

Overall, existing work establishes strong foundations for medical MLLMs, breast imaging models, pathology foundation models, and efficient token compression, but these directions remain largely separate. BreastGPT integrates them into a unified breast cancer workflow model, and BreastStage-Bench evaluates whether such integration can support stage-aware reasoning across heterogeneous modalities and clinical roles.

C BreastStage Curation Details

C.1 Data Sources and Modality Coverage

As described in Section 2 of the main paper, BreastStage integrates 17 sub-datasets across 5 imaging modalities. The goal of this curation is not merely to aggregate public resources, but to reorganize heterogeneous imaging evidence around the clinical progression of breast cancer care: screening, diagnosis, and treatment. This section details the source datasets, selection criteria, quality control procedures, and annotation protocols used to construct each modality subset.

CT Imaging. Computed tomography volumes are sourced from the CT-RATE dataset (Hamamci et al., 2026), which contains 25,692 non-contrast chest CT volumes from 21,304 unique patients. Since breast cancer predominantly affects women, we first filter the dataset to retain only female patients (sex metadata field equal to "F"). We then deploy Qwen2.5-VL-72B (Bai et al., 2025b) as an automated quality inspector, assigning confidence scores (0–10 scale) based on image quality, anatomical coverage, and the presence of breast tissue in the field of view; volumes scoring below the threshold are excluded. For breast cancer screening and diagnosis, we further pass each volume through DRT-M3D (Zhou et al., 2026) to generate automated tumor segmentations and risk assessments, identifying suspicious breast lesions, calcifications, and lymph node involvement visible in the chest CT field of view. After quality filtering and clinical relevance screening, the CT subset of BreastStage contains 20,546 reconstructed volumes from female patients (Table 5), covering both screening-stage opportunistic detection and diagnosis-stage tumor characterization tasks.

Breast Ultrasound (BUS). BUS images are sourced from the **BUS-CoT** dataset (Yu et al., 2026), the first chain-of-thought reasoning breast ultrasound dataset that covers *all* 99 histopathology categories defined by the WHO classification of breast tumors. BUS-CoT contains 11,439 images of 10,019 lesions from 4,838 patients across 18 different ultrasound device types, including B-mode US, Doppler US, and Elastography. Images are collected from open-access publications, publicly available case studies (Radiopaedia, PubMed), and open-access biopsy-confirmed repositories, with each case annotated through a rigorous five-level protocol by six senior breast ultrasound radiologists with 8–26 years of clinical experience: (1) *Observation* (lesion presence, calcification presence), (2) *Feature* (boundary, edge, echo characteristics, calcification feature), (3) *Diagnosis* (BI-RADS score), (4) *Pathology* (benign/malignant + WHO histopathology subtype), and (5) *Chain-of-Thought reasoning* explicitly linking imaging features to the final pathological label. Table 3 summarizes the dataset's pathology distribution. After quality filtering and mapping to our clinical workflow taxonomy (screening-stage BI-RADS triage and diagnosis-stage lesion characterization), the BUS subset of BreastStage contains 10,405 unique image files and 190,730 instruction-following pairs (Table 5), with expert-verified BI-RADS scores, lesion bounding boxes, and chain-of-thought rationales.

	Benign (4,856)	Malignant (4,814)	Others (349)
Categories	Fibroadenoma (1,047)	Invasive ductal carcinoma (896)	Others (349)
	Phyllodes tumour (74)	Invasive lobular carcinoma (75)	
	Intraductal papilloma (72)	Ductal carcinoma in situ (72)	
	Atypical ductal hyperplasia (62)	Mucinous carcinoma (64)	
	Radial scar (56)	Paget disease of the breast (57)	
	Other benign (3,545)	Other malignant (3,650)	

Table 3: Pathology distribution of the BUS-CoT source dataset (Yu et al., 2026), used as the BUS data source for BreastStage. Numbers in parentheses indicate per-subtype lesion counts.

Mammography. Digital mammography is sourced from a subset of **MammoVQA** (Zhu et al., 2025), which originally unifies 15 public mammogram datasets. We adopt the 11 datasets that are licensed for derivative works and have stage-aware annotations compatible with our taxonomy (BMCD, CBIS-DDSM, CDD-CESM, CSAW-M, DMID, EMBED, INbreast, KAU-BCMD, MIAS, RSNA, VinDr-Mammo); we then re-organise their labels around the screening / diagnosis / treatment continuum and add new task templates on top of MammoVQA. Table 4 lists every adopted dataset together with its sample count and the task templates it supports. The union of these 11 sources covers the 9 clinically validated MammoVQA task templates—View (CC/MLO), Laterality, BI-RADS, Pathology, Masking potential, Background tissue, Subtlety, Density, and Abnormality detection (calcification, architectural distortion, mass, asymmetry, etc.).

For BreastStage we use these 11 datasets with patient-level stratification preserved across train and test splits. After mapping the original classification labels to our screening / diagnosis / treatment stage taxonomy and synthesising stage-specific QA pairs, the mammography subset comprises 592,470 unique 2D images and 453,956 instruction-following pairs.

MRI. Multiparametric MRI sequences (T1-weighted, T2-weighted, DWI, DCE-MRI / T1dyn) are acquired from two collaborating clinical institutions [anonymous for review] under IRB approval. Board-certified breast radiologists at the contributing institutions provide expert annotations for tumor segmentation, molecular subtype prediction, and neoadjuvant chemotherapy response assessment. The MRI subset comprises 36,124 unique multiparametric sequence volumes (T1, T1dyn, T2W, DWI, ADC) across screening, diagnosis, and treatment-stage tasks, yielding 926,521 instruction-following pairs.

Whole Slide Imaging (WSI). Histopathology WSIs are sourced from three public repositories—**BCNB**, **TCGA-BRCA**, and **TCGA-HISTAI**—covering invasive ductal carcinoma, invasive lobular carcinoma, ductal carcinoma in situ, and benign / atypical lesions across histological grades I–III. Slides are scanned at 20× or 40× magnification with native resolutions ranging from 80,000×80,000 to 120,000×120,000 pixels. The WSI subset contains 2,510 slides with pathologist-verified molecular subtype labels (ER/PR/HER2 status, Ki-67 index, Nottingham grade), yielding 30,073 instruction-following pairs (captioning, closed- and open-ended VQA, plus prognosis / surgical-planning / systemic-therapy treatment-stage tasks).

Aggregate Statistics. Table 5 summarizes the per-modality data sources adopted by BreastStage. These modality-level counts correspond to the curated image pools before conversion into instruction-following examples. The final instruction corpus expands these image pools into multiple task-specific QA pairs, captions, grounding targets, and report-generation instances according to the stage-aware taxonomy described below.

C.2 Instruction Data Generation Pipeline

This section expands the four-stage curation pipeline introduced in Section 2 of the main paper (*Workflow Data Generation → Stage Instruction Data Construction → Textual Description Generation → Data Splitting and Verification*, illustrated in Figure 2). The four stages are powered by a Data Orchestration Engine that uses two LLMs in complementary roles: **Qwen2.5-VL-72B** (Bai et al., 2025b) is used wherever a decision requires looking at the image (modality-specific quality control, visual-attribute extraction), and **Qwen3-Max** is used for every text-only transformation (Chinese-to-English report parsing, structured-record-to-open-ended-VQA rewriting,

Dataset	Type	Size	Supported Tasks
BMCD (Loizidou et al., 2021)	breast	400	Density, BI-RADS, Laterality
CBIS-DDSM (Lee et al., 2017)	breast	3,103	Density, BI-RADS, Laterality, View
	finding	6,464	Abnormality, Pathology, Subtlety
CDD-CESM (Khaled et al., 2022)	breast	1,003	Density, BI-RADS, Laterality, Pathology
DMID (Oza et al., 2024)	breast	510	Abnormality, Background tissue, Laterality, Pathology, View
	finding	868	Abnormality, Pathology
INbreast (Moreira et al., 2012)	breast	410	Density, Abnormality, BI-RADS, Laterality, View
MIAS (Suckling, 1994)	breast	322	Abnormality, Background tissue, Pathology
	finding	234	Abnormality, Pathology
CSAW-M (Sorkhei et al., 2021)	breast	10,020	Laterality, Masking potential
KAU-BCMD (Alsolami et al., 2021)	breast	2,370	BI-RADS
VinDr-Mammo (Nguyen et al., 2023)	breast	20,000	Density, Abnormality, BI-RADS, Laterality, View
	finding	4,505	Abnormality
RSNA (Carr et al., 2022)	breast	54,705	Density, BI-RADS, Laterality, View
EMBED (Jeong et al., 2023)	exam	72,518	Density (exam), BI-RADS (exam)
Internal total	—	**177,432**	—

Table 4: Internal mammography source datasets aggregated by MammoVQA (Zhu et al., 2025) and adopted as the BreastStage mammography source. "Type" follows the MammoVQA labelling: *breast* = per-image labels, *finding* = per-bounding-box labels, *exam* = per-examination labels.

Modality	Primary Source	Selection Criterion	Unique 2D / 3D Images	Instruction Pairs
CT	CT-RATE (Hamamci et al., 2026)	Female-breast volumes, post-QC	20,546 volumes	254,041
BUS	BUS-CoT (Yu et al., 2026)	All 99 histopathology categories	10,405 images	190,730
Mammo	MammoVQA-derived 11 mammography datasets (Zhu et al., 2025)	Patient-level stratified reorganisation with additional workflow-aware tasks	592,470 images	454,590
MRI	In-house multi-institution multiparametric MRI	Multiparametric studies with complete sequences	36,124 volumes	926,634
WSI	BCNB / TCGA-BRCA / TCGA-HISTAI	Subtype-labeled 20–40× slides	2,510 slides	30,073
Total	**17 sub-datasets**	—	**662,055 images**	**1,856,068**

Table 5: Summary of BreastStage data sources by modality. "Unique 2D / 3D images" counts the distinct image (2D radiology, mammography, histopathology slide) and 3D-volume (CT, MRI multiparametric sequence) files used by the model after filtering each modality to its retained sub-datasets. Each unique image typically backs many QA pairs.

ground-caption synthesis, report generation). Both LLMs are queried with strict JSON output schemas so downstream stages consume structured outputs without bespoke parsing. Per-modality QC prompts and the Chinese-to-English parsers appear in Figure 8, and the stage-aware persona library together with the system-prompt assembly rule appear in Figure 9. The remainder of this section walks through the four stages in one-to-one correspondence with the main-paper subsections.

Workflow Data Generation. The cohort consists of public BUS, mammography, histopathology, and CT data plus an in-house multi-institution MRI dataset acquired from two collaborating hospitals [anonymous for review] under IRB approval; identifying information is removed at the source institutions before any data leaves the hospitals (Appendix A.2). On top of the raw studies, this stage produces three artefacts that the rest of the pipeline relies on: per-image quality scores, lesion-level bounding-box / mask annotations, and stage-aware task labels.

Modality-specific quality control. The visual specialist agents from Section 2 are implemented as Qwen2.5-VL-72B prompted with modality-specific quality-control rules (Fig. 8). The BUS selector rejects convex / sector-probe images outside the breast-ultrasound depth range; the mammography selector strictly rejects chest X-rays and applies PGMI positioning checks plus Eklund implant-displaced views; the breast MRI selector requires bilateral-coil framing or sagittal / axial single-breast views and verifies fat-saturation; the CT

selector requires breast-inclusive FOV. Histopathology WSIs are sourced from a curated public collection where modality and magnification are already standardised, so they bypass this LLM-driven QC step. The selector emits a JSON {validity, reason}; samples flagged as low-quality are not silently dropped but routed to a *Low-Image Check* step in which a breast specialist reviews the image and the LLM's stated reason, and either confirms the rejection or restores the sample.

Bounding-box and mask generation. The four imaging modalities (BUS, mammography, CT, MRI) carry lesion-level spatial annotations, with the source differing by modality. BUS uses the expert hand-drawn lesion masks released with the BUS-CoT dataset (Yu et al., 2026). Mammography (EMBED subset) ships with official lesion bounding boxes; we re-format the released [ymin, xmin, ymax, xmax] tuples into the [xmin, ymin, xmax, ymax] convention used throughout BreastStage, and no mask step is performed on this modality. CT lesion masks come from DRT-M3D (Zhou et al., 2026), which also emits a per-volume cancer-risk score that we retain in the record metadata for the screening-stage CT tasks. MRI lesion masks are drawn manually on T1 and T1dyn sequences by 10 board-certified breast specialists at the contributing hospitals. For the mask-based modalities (BUS, CT, MRI), bounding boxes are derived by connected-component analysis: components with fewer than 50 voxels are discarded as noise, and the tight axis-aligned envelope of every retained component is emitted as a 4-tuple [xmin, ymin, xmax, ymax] for BUS or a 6-tuple [xmin, ymin, zmin, xmax, ymax, zmax] for the 3D volumes (CT, MRI). Histopathology WSIs do not carry pixel-level masks or bounding boxes in BreastStage; they are routed only to the closed/open VQA and report-generation tasks, not to the ground-caption task.

Expert-driven sub-task labelling. After QC and mask generation, the per-dataset clinical tables and reports are reviewed by breast specialists who tag each record with one of **136 sub-task labels** reflecting BI-RADS items, mass / calcification descriptors, biomarker fields, and treatment-stage decisions. Each label is a task template: a high-level sub-task such as "mass margin description" or "HER2 status", typically spanning several leaf fields of the structured report defined in the next stage. Per-modality counts after this stage are reported in Appendix Table 5, and the resulting 136-task taxonomy in Table 6.

Stage Instruction Data Construction. This stage converts each curated study into a schema-bounded structured record, so that no free-text generation step downstream is allowed to introduce clinical content beyond what this record already encodes. The flow follows the four steps introduced in the main paper. (i) **Report translation.** Chinese radiology and pathology reports are translated to English by Qwen3-Max acting as a radiologist agent, using the bilingual prompts in Figure 8. (ii) **Structured report extraction.** Every translated report is then parsed by the radiologist agent into an expert-designed BI-RADS-aligned structured report whose schema is written per modality by breast specialists following clinical reporting guidelines. For example, the MRI structured report (Figure 7) covers per-side breast-level findings (FGT, BPE, post-surgical changes, lymph nodes, BI-RADS, management) plus per-lesion descriptors for masses, non-mass enhancement, and non-enhancing lesions. (iii) **Question bank derivation.** For each task template from Stage 1, breast specialists then write a closed-ended question that draws its answer from one or more leaf fields of the structured report, with options taken directly from the schema's enum values; this constrains BreastStage to questions that the structured record already answers. The resulting tuples carry the form <Task, Question, options, Answer, Images>. (iv) **Task–Stage mapping.** Specialists finally map each task template to one of the three clinical stages (*screening* / *diagnosis* / *treatment*), which is also the key used for task-key balanced sampling on BreastStage-Bench; the per-stage record counts of the final task taxonomy are reported in Table 6.

From the structured records, BreastStage derives two parallel VQA partitions, both grounded in the same schema-bounded source. The *closed-ended* partition packs the schema's enum values as an option list into the question text, so each tuple takes the form <Task, Question+options, Answer, Images> with the answer being one of the options. The *open-ended* partition takes the same structured tuple, drops the option list, and asks Qwen3-Max to rewrite the categorical answer into a fluent clinical description, yielding <Task, Question, Answer, Images> where the answer is free-form text but the underlying clinical fact is unchanged; the rewriting prompt is reproduced in Fig. 6. Because both partitions are pinned to the same structured record, the LLM cannot fabricate clinical content beyond the schema, and the answer cannot leak through the

question text into a text-only shortcut — a property reflected in the $k{=}1$ collapse of BreastGPT's closed-ended accuracy in Fig. 4a, where the model degrades sharply once the visual evidence is removed.

Textual Description Generation. The same structured records that feed VQA construction are also reused to synthesise the two narrative task formats, ground caption and report generation, with the modality-parameterised prompts in Figure 6. Both formats remain pinned to the structured record, so Qwen3-Max is responsible only for the surface form and never for the underlying clinical fact.

Ground caption. Ground captions are produced only for the four imaging modalities that carry bounding boxes (BUS, mammography, CT, MRI); histopathology, which has no boxes, is excluded. Each caption is generated by binding every visible lesion to its box and weaving the resulting visually grounded phrases into a clinical narrative; the surrounding context drawn from the structured record (BI-RADS, risk level, recommendation) is allowed to appear around the boxes but never inside the lesion phrase, so a reader can recover, from the caption alone, which descriptors apply to which spatial region. The modality-specific vocabularies differ in the obvious places (posterior acoustic features for BUS, ACR density for mammography, FGT, BPE and dynamic kinetic curves for MRI), but the visual-vs-context separation is identical across modalities.

Report generation. The report generator targets all five modalities, with one prompt variant per modality so that the output respects the relevant ACR BI-RADS lexicon. Every report is forced into the same four sections (Findings, Impression, Final Assessment, Management) so that downstream evaluation can score them section by section. Three constraints prevent the failure modes typical of pure-LLM report generators: spatial details that are not in the structured input (laterality, lesion size, clock-face position, distance to nipple) cannot appear, source pathology labels must remain probabilistic in the prose ("findings are suspicious for invasive carcinoma" rather than "this is invasive carcinoma"), and any reference to internal field names or coordinates is forbidden so the report reads as clinical dictation. Drafts then flow into the specialist-audit loop in Stage 4 below, where modality-specific prompts are revised and failing records regenerated whenever a per-task flag rate exceeds threshold.

Data Splitting and Verification. Records are split between training and BreastStage-Bench at patient-id granularity; because the source datasets cover disjoint patient populations, no cross-dataset patient deduplication is required. The split itself uses a single global stratified sampler over the composite (modality, task type, pathology label) key, so the resulting train and test partitions match on the underlying task and class distributions. BreastStage-Bench equals the test partition, drawn with task-key balanced sampling so that every task template in Table 6 is represented. Verification has two layers. Automated heuristics first remove records with malformed bounding-box coordinates, invalid option labels, instruction–answer conflicts, hallucinated clinical terms, and near-duplicate QA pairs (MinHash > 0.85). Three breast specialists then conduct an independent task-level audit on Bench: each specialist reviews 5 random samples per task and flags clinically inconsistent or stage-mismatched cases. Any task whose flag rate is non-negligible across specialists triggers re-running the relevant LLM stage with the failing prompts revised.

Stage-aware system prompts. BreastStage uses two complementary prompt families. The data-orchestration prompts in Figure 8 and Figure 6 drive the curation pipeline: the five orchestration agents (BUS / Mammography / Breast MRI Quality Selectors, Side-Aware Report Splitter, Bilingual Pathology Report Parser) and the three generators (Open-VQA, Grounded Caption, Report). At training and inference time, the resulting instruction-following pairs are then served with stage-specific system prompts that switch BreastGPT between three clinical roles: screening radiologist, diagnostic radiologist, and treatment-stage breast oncologist or pathologist; Figure 9 reproduces the three personas and the deterministic template that combines a persona with a modality-aware task instruction.

C.3 Task Taxonomy and Distribution

We summarise the BreastStage instruction corpus along two complementary axes. Table 6 groups the 136 task templates into a 12-category taxonomy under the three clinical stages, and Figure 10 reports the joint distribution over (clinical stage, imaging

Open-VQA Generator

You are a medical imaging data assistant specialising in breast `{modality}`. Your task is to rewrite a structured attribute into a natural-language Question–Answer pair suitable for a Breast `{modality}` VQA dataset.

Rules.

1. The `Question` *must* be open-ended and descriptive: do *not* ask yes/no questions, do *not* include options or binary choices.
2. The `Answer` *must* be a complete sentence that is strictly based on the provided `Answer` value, and does *not* introduce any new medical interpretation, diagnosis, or inference.
3. Do *not* introduce information not explicitly contained in the input.
4. Do *not* repeat or list the provided options in the output.
5. Use varied sentence structures across different samples.
6. Use precise anatomical and imaging terminology appropriate for breast `{modality}`.
7. Output *only* valid JSON in the specified format. No extra text.

Input (JSON): `{"Question": "...", "options": ["...","...","..."], "Answer": "..."}`.
Output (JSON): `{"Question": "...", "Answer": "..."}`.

Grounded Caption Generator

You are an **expert Radiologist Assistant** producing a **Grounded Caption** for breast `{modality}`. Strictly separate *visual attributes* (shape, margin, echo / signal / density, posterior or kinetic features) from *clinical context* (BI-RADS, risk, diagnosis, management).

Step 1 – Reference phrase. For each finding, write a noun phrase that names *only what is visible inside the bounding box*: include side, view (when applicable), abnormality type, and the relevant visual descriptors. Exclude BI-RADS, risk, diagnosis, and management.

Step 2 – Caption skeleton. Compose a professional radiology sentence with placeholder tokens `<ref-object><bbox>` immediately after each lesion phrase; inject the excluded clinical context (BI-RADS, risk, recommendation) *around* the placeholders, never inside the reference phrase.

Step 3 – Strict alignment. Emit `ref[]`, `bbox[]` as flat lists in caption order; `ref[i]` corresponds to `bbox[i]` and to the i-th `<ref-object><bbox>` token. Boxes are 2D `[xmin,ymin,xmax,ymax]` for BUS / mammography / histopathology and 3D `[xmin,ymin,zmin,xmax,ymax,zmax]` for CT / MRI.

Output JSON: `{caption, objects:{ref, bbox}}`.

Report Generator

You are a **board-certified breast radiologist**. Generate a formal breast `{modality}` report from the structured `patient_data`, compliant with the **ACR BI-RADS Lexicon** for that modality and reflecting real-world dictation style.

Hard safety rules. (1) No fabrication of laterality, lesion size, clock-face position, or distance to nipple unless explicitly present in `patient_data`; if missing, refer generically ("the lesion"). (2) No histological / molecular confirmation: replace "X is invasive carcinoma" with probabilistic phrasing ("findings are suspicious for X"). (3) Bounding boxes are internal; never mention coordinates, "bbox", or any data-structure terminology in the prose.

Sections (fixed order). *Findings*: one passive-voice paragraph per lesion, ordered by descending BI-RADS, integrating only descriptors present in `patient_data`. *Impression*: probabilistic synthesis. *Final Assessment*: the highest BI-RADS category among findings. *Management*: ACR-aligned recommendation (Cat 2 – routine; Cat 3 – short-interval follow-up; Cat 4–5 – biopsy / tissue diagnosis). Surgical / systemic therapy decisions are out of scope.

Output: four-section markdown with the section names above.

Figure 6: Generator prompts that produce the BreastStage instruction-following partitions from each structured record: the open-ended VQA generator (top), the grounded caption generator (middle), and the report generator (bottom). All three are instantiated per modality with `{modality}` $\in$ {BUS, mammography, CT, MRI, histopathology}.

modality, task category) that this taxonomy induces.

C.4 Human Annotator Statement

All clinical annotators are full-time staff at the collaborating institutions and contributed as part of an ongoing research collaboration under a signed institutional agreement, rather than via crowdsourcing or per-item piecework. Each annotator was briefed in writing on the labeling rubric, the upstream de-identification guarantees, and the intended research use, and gave written consent for their annotations to be released with BreastStage.

D BreastGPT Technical Details

D.1 Base Model Configuration

BreastGPT is initialized from Qwen3-VL-8B-Instruct and inherits its instruction-following multimodal interface. We keep the language backbone and the standard radiology visual encoder largely aligned with the released model and introduce additional components only where the breast cancer workflow requires capabilities not naturally covered by the base model. The major architectural extension is the GigaPixel branch for WSI processing together with the coverage-maximizing token selector that maps heterogeneous image representations into a fixed downstream token budget. For CT,

MRI Structured Report Template

```
{
  "breast_level": {
    "FGT": {"left": "C", "right": "C"}, // "A","B","C","D"
    "BPE": {"left": "Moderate", "right": "Mild"}, // "Minimal","Mild","Moderate","Marked"
    "breast_symmetry": true,
    "post_surgical_changes": {
      "left": {"scar_tissue": false, "lumpectomy_changes": false, "mastectomy_changes": false},
      "right": {"scar_tissue": false, "lumpectomy_changes": false, "mastectomy_changes": false}
    },
    "associated_features": {
      "left": {
        "nipple_retraction": null, "nipple_invasion": null,
        "skin_retraction": null, "skin_thickening": null,
        "skin_invasion": {"direct_invasion": null, "inflammatory_cancer": null},
        "axillary_adenopathy": null, "pectoralis_muscle_invasion": null,
        "chest_wall_invasion": null, "architectural_distortion": null,
        "ductal_dilation": null // "mild","marked",null
      },
      "right": { ... same schema as "left" ... }
    },
    "lymph_nodes": {
      "axillary_left": null, "axillary_right": null, // "normal","abnormal"
      "internal_mammary_left": null, "internal_mammary_right": null
    },
    "BI-RADS": {"left": null, "right": null}, // "0","1","2","3","4A","4B","4C","5","6"
    "management":{"follow_up_recommended": null, "biopsy_recommended": null}
  },
  "mass": [
    {
      "id": 1, "side": "left",
      "location": {"quadrant": "UOQ", // "UIQ","UOQ","LIQ","LOQ","central",null
                   "depth": "middle"}, // "anterior","middle","posterior",null
      "size_mm": {"long": 12, "short": 9, "depth": 8},
      "signal_characteristics": {
        "T1W": "hyperintense", "T2W": "hyperintense", // "hyperintense","isointense","hypointense",null
        "DWI_restriction": "present", // "present","absent",null
        "ADC": "hyperintense",
        "kinetics": {"initial": "Fast", // "Slow","Medium","Fast",null
                     "delayed": "Wash-out"} // "Persistent","Plateau","Wash-out",null
      },
      "morphology": {"shape": "oval", // "round","oval","lobular","irregular",null
                     "margin": "circumscribed"}, // "circumscribed","irregular","spiculated",null
      "internal_enhancement": {
        "pattern": "heterogeneous", // "homogeneous","heterogeneous","rim_enhancement","dark_internal_septations",null
        "associated_vascularity": {"adjacent_vessel_sign": null,
                                   "increased_peritumoral_vascularity": null}
      }
    }
  ],
  "non_mass_enhancement": [
    {
      "lesion_id": 1, "side": "left",
      "location": {"quadrant": null, "depth": null},
      "size_mm": {"long": null, "short": null, "depth": null},
      "distribution": "Segmental", // "Focal","Linear","Regional","Segmental","Multiple regions","Diffuse",null
      "internal_enhancement_pattern": "Clustered_ring", // "Homogeneous","Heterogeneous","Clustered","Clustered_ring",null
      "signal_characteristics": { ... same as "mass" ... }
    }
  ],
  "non_enhancing_lesion": [
    {
      "lesion_id": 1, "side": "left",
      "location": {"quadrant": null, "depth": null},
      "size_mm": {"long": null, "short": null, "depth": null},
      "signal_characteristics": { ... same as "mass" ... },
      "lesion_type_detail": "ductal_precontrast_high_signal_on_T1W"
                            // "cyst","non_enhancing_mass","architectural_distortion",null
    }
  ]
}
```

Figure 7: Expert-designed BI-RADS-aligned structured template used for the MRI cohort. Every leaf field is either a numeric measurement, a Boolean flag, or one of a fixed enum list (shown as a comment). Reports are first translated to English by Qwen3-Max and then populated into this template under its enum and type constraints; specialists then derive a question per `task template` whose answer is drawn from one or more leaf fields. Schema-conformant population guarantees that the resulting QA pairs cannot introduce facts that the source report does not already record.

BUS, mammography, and MRI inputs, the standard branch is used; for WSI inputs, modality-aware routing activates the GigaPixel branch described below.

D.2 GigaPixel Encoder Implementation

The GigaPixel encoder is designed to preserve diagnostically relevant pathology patterns without forcing the LLM to attend over the full patch sequence. A single WSI can contain tens or hundreds of thou-

BUS Quality Selector

You are a **Breast Ultrasound Quality Specialist**. Your goal is to validate if the image is likely a **Breast/Axilla Ultrasound**. Evaluate the following:

1. Modality & Probe Geometry (BREAST FOCUS):

- *Linear Probe (Rectangular Field):* high probability of valid Breast US.
- *Convex/Sector Probe (Fan-shaped Field):* **WARNING** – typical for Abdomen/Liver. Mark Valid only if the text explicitly confirms a deep breast mass.
- *Modality:* confirm Ultrasound (speckle texture).

2. Image Quality & Settings:

- *Depth:* $<$ 4–6 cm; reject if deep abdominal structures (liver / kidney) are visible.
- *Focus zone:* aligned with the mammary layer.
- *Gain:* fat / glandular interface distinguishable.

3. Text–Image Alignment: flag if the text mentions "Liver", "Gallbladder", or "Thyroid".

Output JSON: {validity: Valid/Invalid, reason}.

Mammography Quality Selector

You are a **Mammography Quality Control Expert**. Validate if the image is a diagnostic **Mammogram (MG)**. Evaluate the following:

1. Modality Distinction (CRITICAL):

- Confirm Mammogram: soft-tissue breast against black air background.
- **STRICTLY REJECT** Chest X-Ray (CXR): if you see ribs, lungs, clavicles, or the spine, mark Invalid.

2. Breast-Specific Positioning (PGMI Criteria):

- *Pectoralis muscle:* for MLO views, must extend to the nipple line.
- *Nipple profile:* in profile (not superimposed) unless retracted.
- *Skin folds:* reject severe folds that mimic abnormalities.
- *Implants:* identify "Implant Displaced (Eklund)" views vs. standard views.

3. Exposure & Contrast: adipose (dark gray) vs. fibroglandular tissue (white) must be distinguishable; reject motion blur at calcifications.

4. Text–Image Alignment: text describes mammographic findings (asymmetry, distortion, calcification).

Output JSON: {validity, reason}.

Breast MRI Quality Selector

You are a **Breast MRI Physics Specialist**. Validate if the image is a diagnostic **Breast MRI** sequence. Evaluate the following:

1. Coil & Field of View (BREAST FOCUS):

- *Breast coil geometry:* bilateral breasts hanging in the coil, or a dedicated sagittal/axial single-breast view.
- *Exclude:* Brain MRI, Spine MRI, or Abdominal MRI where the breast is incidental and compressed.
- *FOV:* must include the axillary tail region.

2. Sequence & Physics:

- *Fat Saturation:* **CRITICAL** for breast – fat must be suppressed (dark) except in T1 non-fat-sat anatomy scans.
- *Cardiac motion artefacts:* check the left-breast phase-encoding direction; severe pulsation rendering the medial breast unreadable $\rightarrow$ Invalid.
- *Silicone select:* for implant cases, confirm silicone-specific sequences are used.

3. Text–Image Alignment: text describes "enhancement kinetics", "BPE", or "washout"; flag mismatch if it describes CT density.

Output JSON: {validity, reason}.

Side-Aware Report Splitter

You are an **expert breast radiologist**. Process Chinese radiology reports: (1) split findings by side, (2) translate everything into professional medical English, (3) determine benign / malignant from findings and BI-RADS.

CRITICAL RULE – NO CHINESE. The final JSON output must not contain any Chinese characters; the thinking and the result are 100% English.

Side-assignment rules (per sentence):

- left-side keyword (*zuo / zuoce / zuoru*) $\rightarrow$ LEFT.
- right-side keyword (*you / youce / youru*) $\rightarrow$ RIGHT.
- bilateral / no side mentioned $\rightarrow$ include in BOTH, rewritten per side.
- Sentences mentioning both sides (e.g. "Left is smaller than right") are split into two side-specific English statements.

Translation standard: use professional English medical terminology ("hyperplasia", "BI-RADS", "outer upper quadrant"); preserve every measurement, echo, and vascularity descriptor.

Output JSON: {left: {finding, ...}, right: {finding, ...}}.

Bilingual Pathology Report Parser

You are an **expert Pathologist and Data Structuring Specialist**. Extract information from raw Chinese breast pathology reports and convert it into a structured **English JSON** format.

Input format. {`type`: "Surgical pathology" or "Non-surgical pathology"; `pathology`: report text}. Surgical pathology is text fields separated by spaces or colons. Non-surgical pathology is often a Python set string such as `"{'text', 'text'}"`.

Symbol decoding. `(-)` $\rightarrow$ Negative; `(+)` $\rightarrow$ Positive; `/` $\rightarrow$ Not Applicable or Not Found.

Translation rules (strict adherence). Histology terms map deterministically: *jin run xing ai* $\rightarrow$ Invasive Carcinoma; *fei te shu lei xing* (NST) $\rightarrow$ Invasive Carcinoma of No Special Type (NST); *dao guan yuan wei ai* $\rightarrow$ Ductal Carcinoma In Situ (DCIS); *xian wei xian liu* $\rightarrow$ Fibroadenoma; *ying hua xing xian bing* $\rightarrow$ Sclerosing Adenosis; *ru xian bing* $\rightarrow$ Mastopathy; *suo xing xi bao* $\rightarrow$ Spindle Cell; ... Biomarkers (ER, PR, HER2, Ki-67, Nottingham grade) and molecular subtype are normalized to the canonical Enum (Luminal A / B, HER2-enriched, TNBC).

Output JSON: canonical pathology schema; missing fields left null.

Figure 8: System prompts of the five BreastStage data-orchestration agents (modality-specific quality selectors plus the bilingual report splitter and pathology parser). Each agent runs Qwen3-Max with a strict JSON output schema, so its decision can be consumed by the next stage without manual parsing. Generator prompts that produce the actual instruction-following pairs are reproduced in Figure 6.

Persona Library; one of these three blocks is selected by the record's `stage` *field:*

Screening Persona

You are an expert screening breast radiologist. Your clinical focus is on early detection and risk triage across broad populations. You possess a high sensitivity for identifying subtle abnormalities within varied anatomical backgrounds to determine if further investigation is warranted.

Used for: `stage = screen` (mammography, BUS, opportunistic CT).

Diagnosis Persona

You are an expert diagnostic breast radiologist. Your clinical focus is on definitive lesion characterization and differential diagnosis. You rigorously analyse imaging features to separate benign from malignant entities, providing conclusive assessments that directly guide patient management.

Used for: `stage = diagnosis` (BUS, mammography, multiparametric MRI).

Treatment Persona

You are an expert breast oncologist and pathologist. Your clinical focus is on decoding tumour biology and evaluating disease progression. You synthesise pathological and imaging data to assess therapeutic responses and architect comprehensive treatment or surgical strategies.

Used for: `stage = treatment` (treatment-stage MRI, histopathology WSI).

⇓ *combined into the final system prompt by the rule below* ⇓

Assembly Rule (combines the persona above with a task instruction)

The final system prompt is a deterministic concatenation of three components:

`persona` + "Now, apply your expertise to the following task:" + `instruction(task, modality)`

Persona slot. Filled from the persona library above according to the record's `stage` field (*screen* / *diagnosis* / *treatment*); for stage-less records a generic breast-clinician persona is used.

Instruction slot. Filled from one of five task templates, each carrying a {modality} placeholder:

- *closed_vqa:* "Answer the multiple-choice question based on the provided {modality}. Respond ONLY with the exact option chosen."
- *open_vqa:* "Answer the open-ended question based on the provided {modality}."
- *caption:* "Generate a concise descriptive caption for the provided {modality}."
- *groundcaption:* "Generate a descriptive caption for the provided {modality}. You MUST explicitly ground each mentioned finding with a {2D|3D} bounding box."
- *report:* "Generate a comprehensive clinical report based on the provided {modality}."

Modality and bounding-box format. The {modality} placeholder takes the canonical noun phrase per modality: BUS / Mammo / Histopathology use 2D images and 2D bounding boxes; CT and MRI use 3D volumes and 3D bounding boxes.

Figure 9: Stage-aware system-prompt library used by BreastGPT. The three personas (top) cover the screening / diagnosis / treatment workflow; the assembly rule (bottom) deterministically composes a persona, a transition phrase, and a modality-specific task instruction at training and inference time. The same persona is reused regardless of modality, so a single model can adopt stage-appropriate reasoning style without task-specific heads.

sands of tiles, and many clinically important cues are sparse: mitotic figures, ductal structures, invasive fronts, necrosis, lymphocytic infiltration, and tumor-stroma interfaces may occupy only a small fraction of the slide. The WSI branch therefore separates representation learning into three steps: local patch encoding, long-range contextualization, and query-aware coverage selection.

Stage 1: Patch Extraction and Encoding. WSI tiles are extracted at 20× magnification with non-overlapping 512×512 pixel patches. For a typical breast WSI this yields tens of thousands of patches. Each patch is encoded once using the frozen CONCH v1.5 foundation model (Lu et al., 2024), producing 512-dimensional embeddings that are cached to disk for reuse:

$$\mathbf{v}_i = \text{CONCH}(\mathbf{x}_i), \quad i = 1, \ldots, N, \quad \mathbf{v}_i \in \mathbb{R}^{512}.$$

Stage 2: LongNet Contextualization. The patch embeddings are processed by a 2-layer LongNet encoder with dilated attention (matching the configuration described in Section 3 of the main paper). The dilation rates follow an exponential schedule, allowing the model to capture both local cellular details and global tissue architecture. The contextualized embeddings are:

$$\mathbf{h}_i = \text{LongNet}(\{\mathbf{v}_1, \ldots, \mathbf{v}_N\})_i.$$

LongNet's exponentially growing dilation rates give the encoder long-range coverage without the $O(N^2)$ cost of full attention.

D.3 Coverage-Maximizing Token Selection

The coverage-maximizing token selector is the mechanism that makes BreastGPT practical as a single model across ordinary radiology images and gigapixel pathology slides. Without compression, WSI inputs would exceed the LLM context budget by several orders of magnitude; with naive pooling or truncation, sparse but clinically important

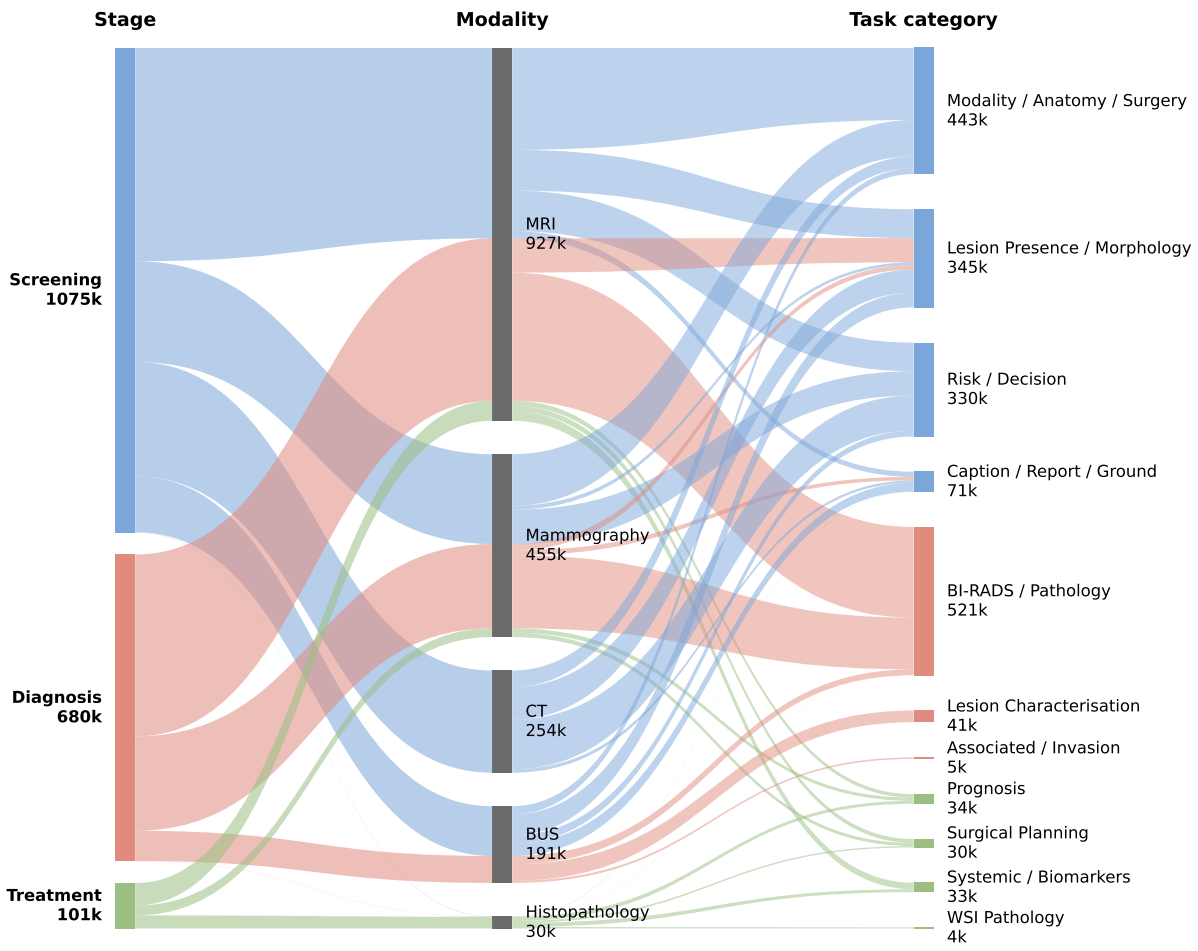


Figure 10: Joint distribution of the 1.86M BreastStage instruction-following pairs across clinical stage (left), imaging modality (middle), and task category (right). Ribbon width is proportional to record count, and ribbons are coloured by clinical stage so each contributing stage remains identifiable on both sides of the modality node.

regions may be removed. We therefore use a coverage objective that keeps the selected tokens both *clinically query-relevant* and *globally representative* of the visual evidence.

Given a sequence of projected visual tokens and an associated clinical instruction, the goal of the selector is to identify a compact subset $S \subset \{1, \ldots, N\}$ with $|S| = k$ ($k = 128$ by default) that jointly preserves relevance to the clinical query and fidelity to the overall visual distribution. For WSI inputs, the source sequence is the LongNet-contextualized patch representation $\{h_1, \ldots, h_N\}$; for standard radiological inputs, the source sequence is the ViT token sequence. In both cases, the selected tokens are passed to the LLM as the fixed-budget visual context. This gives BreastGPT a modality-invariant interface: each clinical image, regardless of its original resolution, is represented by the same number of visual tokens before language-model reasoning.

Similarity matrices. Let $\{v_1, \ldots, v_N\}$ denote the projected vision tokens after the LLM projection layer, aligned with the text token space, and let $\{v'_1, \ldots, v'_N\}$ denote the pre-projection tokens that retain modality-native similarity structure. Let $\{t_1, \ldots, t_m\}$ denote the text tokens encoding the clinical instruction. We define text–vision (T-V) and vision–vision (V-V) similarity matrices as

$$M^{tv}_{i,j} = t_i^\top v_j, \quad M^{vv}_{i,j} = v'^\top_i v'_j, \\ \|t_i\|_2 = \|v_j\|_2 = \|v'_j\|_2 = 1. \tag{2}$$

The similarities are then softmax-calibrated with temperatures τ_t and τ_v:

$$\widetilde{M}^{tv}_{i,j} = \frac{\exp(M^{tv}_{i,j}/\tau_t)}{\sum_{j'} \exp(M^{tv}_{i,j'}/\tau_t)}, \\ \widetilde{M}^{vv}_{i,j} = \frac{\exp(M^{vv}_{i,j}/\tau_v)}{\sum_{j'} \exp(M^{vv}_{i,j'}/\tau_v)}. \tag{3}$$

Clinical Stage	Task Category	Representative Task Templates	Samples
Screening (1,075,092)	Modality / Anatomy / Surgical History	Modality, View, Laterality, Orientation, ACR, FGT, BPE, Background tissue, breast_symmetry, post_surgical_changes (mastectomy / lumpectomy / scar)	442,860
	Lesion Presence & Morphology	abnormal_presence, abnormality_type, mass, non_mass_enhancement, non_enhancing_lesion, Shape, LesionEdge, LesionBoundary, EchoCharacteristics, LesionCalcificationFeatures, BloodFlow, lesion depth	244,188
	Risk Assessment & Screening Decision	RiskLevel, ScreeningDecision, Density_Risk, follow_up_recommended	329,756
	Captioning / Report / Grounding	ground caption (BUS, CT), report generation (BUS), Report (MRI)	58,288
Diagnosis (680,409)	BI-RADS & Pathology Diagnosis	BI-RADS, birads, diagnosis, Pathology, main_type	520,969
	Lesion Characterization	mass.morphology (margin, shape), mass.signal (T1W, T2W, ADC, DWI, kinetics), internal_enhancement, calcificationType, boundaryDefinition, echoTexture, morphology, distribution	141,761
	Captioning / Report / Grounding (Mammo / MRI)	ground caption (Mammo), report generation (Mammography, MRI)	12,227
	Associated Findings & Invasion	lymph_nodes (axillary, internal_mammary), skin_thickening, nipple_retraction, architectural_distortion, ductal_dilation, pectoralis / chest_wall / nipple / skin invasion, vascularity, invasiveness	5,452
Treatment (100,567)	Prognosis & Outcome Prediction	prognosis, Treatment.Prognosis, Treatment.treatment.prognosis	33,743
	Surgical Planning & Urgency	surgical_plan, SurgicalPlan, treatment.urgency	29,797
	Systemic Therapy & Molecular Biomarkers	systemic_therapy, biomarkers (ER, PR, HER2, Ki67, molecular_subtype)	33,200
	WSI Pathology Subtyping & Captioning	Tumor, ER, PR, HER2, HER2 Expression, Molecular subtype, Histological grading, slide caption / VQA	3,827

Table 6: Twelve-category task taxonomy of BreastStage. Templates are grouped by clinical stage (Screening, Diagnosis, Treatment) and, within each stage, by the family of clinical concept being assessed. The *Samples* column reports the number of instruction-following pairs whose majority clinical stage falls in that category.

Dual coverage objective. We select S by maximizing the joint coverage score

$$f(S) = \underbrace{\frac{1}{m}\sum_{i=1}^{m}\max_{s\in S}\widetilde{M}^{tv}_{i,s}}_{\text{T-V coverage}} + \alpha\cdot\underbrace{\frac{1}{N}\sum_{i=1}^{N}\max_{s\in S}\widetilde{M}^{vv}_{i,s}}_{\text{V-V coverage}}, \tag{4}$$

where α balances query relevance against global visual representativeness. The T-V term rewards tokens that explain the clinical instruction, while the V-V term forces the selected subset to cover the full visual distribution so that diagnostically relevant but less immediately salient regions are not discarded.

Why both terms are needed. The two coverage terms play complementary roles. T-V coverage favors tokens aligned with the prompt, such as lesion regions for BI-RADS questions or tumor-cell regions for molecular subtype questions. However, relying on T-V coverage alone can over-concentrate the selected subset around a small number of highly salient regions. V-V coverage counteracts this collapse by selecting tokens that represent the broader slide or image distribution. This is especially important for pathology, where grade, subtype, and treatment-response cues can depend on tissue architecture, tumor-stroma interaction, and heterogeneity across multiple regions. Conversely, V-V coverage alone may preserve visually diverse background regions that are irrelevant to the current clinical task. The combined objective gives BreastGPT a principled way to balance prompt alignment and visual diversity.

Greedy optimization. Each component of f is a facility-location function over the token set and is therefore monotone submodular; the non-negative

linear combination remains monotone submodular. A standard greedy algorithm that iteratively adds the token with the largest marginal gain achieves the usual $(1 - 1/e)$ approximation to the NP-hard optimum in $O(kN)$ time. End-to-end measured selector latency at $k = 128$ on a representative histopathology WSI is reported in Table 15.

Implementation details. The selector is applied after modality-specific visual encoding and before LLM injection. For standard radiology images, token selection operates on the ViT token sequence. For WSI inputs, token selection operates on LongNet-contextualized patch embeddings. The same value of k is used for both branches, which ensures that radiology and pathology inputs have comparable downstream inference costs. During training, the selector is run on-the-fly so that the language model learns to reason over the same type of compact visual context used at inference time.

D.4 Training Configuration

BreastGPT is trained in two stages. In *Stage 1*, we warm up the visual front end on the closed-VQA subset of BreastStage: the LLM backbone is frozen, while the ViT, the aligner (which includes the WSI projection and LongNet module), and the resolution gate are trainable; the CONCH v1.5 encoder is always frozen. This step aligns the visual representations, both the ViT branch for radiology and the CONCH+LongNet branch for WSIs, with the multimodal token space before any gradient flows into the language model. In *Stage 2*, we unfreeze the LLM and continue end-to-end fine-tuning on all four task formats across every modality; the CONCH v1.5 encoder remains frozen. The coverage token selector is training-free and applied identically at training and inference. We use AdamW with a cosine learning rate schedule and DeepSpeed ZeRO-2, cap the per-image visual budget at 1024 tokens, cap CT and MRI volumes at $384 \times 384 \times 48$, and fix the selector budget at $k = 128$ for both branches so that radiological and pathological inputs incur the same downstream inference cost. Table 7 summarises the two-stage training configuration used for BreastGPT, taken directly from the launch scripts.

D.5 System Prompt Templates

The full stage-aware system-prompt library used by BreastGPT, comprising three personas (Screening, Diagnosis, Treatment, plus a generic clinician fallback) and the deterministic rule that composes a persona with a modality-specific task instruction, is reproduced verbatim from the codebase in Figure 9. The same library is used both during training and at inference time; the persona slot is selected by the record's `stage` field, and the instruction slot is filled by task and modality as described in the assembly rule.

E BreastStage-Bench Evaluation

E.1 Benchmark Construction

BreastStage-Bench is constructed to evaluate whether a model can preserve stage-aware clinical reasoning across the full breast cancer workflow rather than merely solve isolated image-classification tasks. We use patient-level stratified sampling and keep all task instances derived from the same patient within the same split. This prevents leakage through repeated images, repeated reports, or semantically equivalent QA pairs.

The benchmark construction follows three principles:

1. **No data leakage**: Strict patient-level separation between training and test sets
2. **Task diversity**: All 136 task templates are represented with balanced class distributions
3. **Clinical realism**: Test cases reflect real-world prevalence and complexity

BreastStage-Bench contains 12,182 evaluation records in total, distributed across the four task families and five modalities as follows:

- Closed-ended VQA (5,369 questions): BUS 987, CT 825, Mammography 1,828, MRI 908, Histopathology 821.
- Open-ended VQA (2,833 questions): BUS 987, CT 825, MRI 908, Histopathology 113.
- Ground caption (2,910 cases): BUS 1,000, CT 510, Mammography 1,000, MRI 400.
- Caption (70 cases): Histopathology slide-level captioning.
- Report generation (1,000 cases): multiparametric MRI structured reports.

E.2 Evaluation Metrics

Closed-ended VQA. We adopt accuracy as the primary metric. For multiple-choice questions, we

Setting	Stage 1: Visual front-end warm-up	Stage 2: End-to-End SFT
Epochs	1	2
Trainable modules	ViT, aligner (incl. WSI projection / LongNet)	LLM, ViT, aligner
Frozen modules	LLM backbone, CONCH encoder	CONCH encoder
Datasets	Closed-VQA subsets (BUS, CT, Mammo, MRI) and the histopathology caption / closed / open VQA tasks	All four task formats (closed, open, grounding, report) on every modality (16 splits)
Initialization	Released Qwen3-VL-8B-Instruct weights	Stage-1 checkpoint
Optimizer	AdamW (default β, weight decay 0.1)	AdamW (default β, weight decay 0.1)
Learning rate	1×10^{-5}, cosine schedule, 0.03 warm-up ratio	1×10^{-5}, cosine schedule, 0.03 warm-up ratio
Gradient clipping	1.0	1.0
Precision	bfloat16	bfloat16
Image / volume budget	`IMAGE_MAX_TOKEN_NUM`=1024; CT and MRI volumes capped at $384 \times 384 \times 48$	same
Max sequence length	default	10,000 tokens
Selector budget k	128	128
Distributed	DeepSpeed ZeRO-2; per-device batch and accumulation set per launch	DeepSpeed ZeRO-2; per-device batch and accumulation set per launch

Table 7: Training configuration for BreastGPT. Both stages use the `ms-swift sft` entry point with FlashAttention, gradient checkpointing, and DeepSpeed ZeRO-2.

extract model responses using robust regular expressions matching option patterns (A/B/C/D), answer text, and common variants such as "Option A" or "the correct answer is A". If no valid option is detected in the first 100 tokens of the response, the prediction is marked invalid and counted as incorrect. This avoids rewarding models for verbose but non-committal answers.

Open-ended VQA. Unlike LLM-as-judge protocols commonly used in recent multimodal evaluation toolkits (Yu et al., 2023, 2024; Duan et al., 2024), BreastStage-Bench reports automatic semantic and lexical scores for open-ended responses, and uses breast-expert sampling validation to audit the clinical correctness of references and representative outputs. For normalized open-ended scores, each evaluated response receives a score $s_i \in [0, 1]$ according to the task-specific scoring rubric, and the total score is:

$$S = \frac{\sum_{i=1}^{N} s_i}{N} \times 100\% \tag{5}$$

The few-shot prompt includes 9 in-context examples covering fully correct (1.0), entirely incorrect (0.0), and various "partially correct" responses (0.25, 0.5, 0.75). The evaluator is instructed to prioritize clinical correctness over surface-form similarity and to penalize unsupported diagnosis, missing contraindications, or recommendations inconsistent with the specified clinical stage.

Grounding Tasks. For lesion localization, we compute Intersection over Union (IoU) between predicted and ground-truth bounding boxes:

$$\text{IoU} = \frac{\text{Area}(\text{pred} \cap \text{gt})}{\text{Area}(\text{pred} \cup \text{gt})} \tag{6}$$

A prediction is considered correct if IoU ≥ 0.5. We report mean IoU and accuracy@0.5 across all grounding tasks. Predictions with malformed coordinate syntax are counted as invalid and assigned IoU 0.

Report Generation. We evaluate generated reports using:

- **BLEU-4**: Measures n-gram overlap with reference reports
- **BERT Score**: Captures semantic similarity using contextual embeddings (microsoft/BiomedNLP-PubMedBERT-base)
- **ROUGE-1**: Assesses recall of unigrams

The weighted composite score is: Wtd $= 0.5 \times$ BERT $+ 0.25 \times$ BLEU $+ 0.25 \times$ ROUGE-1.

E.3 Breast Expert Validation

BreastStage and BreastStage-Bench are constructed with a structured expert sampling audit rather than relying only on automatic filters or LLM-generated text. Three board-certified breast specialists (two breast surgeons and one breast radiologist, each with at least five years of post-residency experience) independently review a stratified subset of BreastStage-Bench, scoring each item along three pre-registered dimensions that mirror the claim made in the main text: *task validity*, *answer correctness*, and *clinical consistency*. The audit covers all five modalities (BUS, mammography, CT, MRI, histopathology) and all task families (closed VQA, open VQA, ground caption, report).

Sampling protocol. Sampling is stratified by (modality, task template) so that every one of the 136 task templates contributes evenly: each specialist independently audits 5 random samples per template, yielding 680 per-specialist reviews and a pooled review pool of 2,040 *(record, specialist)* judgments. Specialists are blind to which LLM stage produced each record. Each item is judged on the following three binary dimensions:

- **Task validity** — the question is well-posed for the indicated clinical stage and the requested evidence is identifiable in the image (or in the source report, for report-only items).
- **Answer correctness** — the gold answer is logically consistent with the question and contains no claims unsupported by the image or source report (no hallucinated lesions, biomarkers, or risk categories).
- **Clinical consistency** — the assigned clinical stage (screening / diagnosis / treatment) and the modality-specific terminology (BI-RADS category, ACR density, molecular subtype, etc.) match the intended clinical role and current breast-oncology practice.

Disagreement resolution. Items flagged by at least one but not all specialists are routed to a *consensus review*: the three specialists meet, the contested record is re-examined together, and a single binary decision is recorded by majority vote (with the breast radiologist breaking ties on imaging-specific calls and the senior breast surgeon breaking ties on management calls). Templates whose post-consensus failure rate exceeds 10 % are routed back to the responsible LLM prompt for revision, after which the affected records are regenerated and re-audited.

We report per-dimension approval rate and inter-rater agreement (Fleiss's κ) in Table 8. Across all three dimensions the post-revision approval rate exceeds 95 %, and Fleiss's κ stays in the substantial-agreement range ($\kappa \in [0.74, 0.86]$), indicating that the audit signal is reliable rather than dominated by a single specialist's idiosyncrasies.

Dimension	Reviews	Pre-rev.	Post-rev.	κ	Re-gen.
Task validity	2,040	91.4%	98.2%	0.78	14
Answer correctness	1,620	90.6%	97.4%	0.79	11
Clinical consistency	2,040	96.6%	99.8%	0.86	2
Overall (any-dim. fail)	2,040	**86.1%**	**96.3%**	0.74	22

Table 8: Expert validation results on BreastStage-Bench. *Pre-rev.* and *Post-rev.* denote approval rates before and after prompt revision. κ is Fleiss's κ among the three specialists on binary approval decisions. *Reviews* counts *(record, specialist)* judgments. *Answer correctness* is evaluated only on tasks with gold answers or free-text generation targets, resulting in a smaller evaluation pool. *Re-gen.* denotes the number of revised task templates.

Per-stage breakdown. Post-revision approval is stable across clinical stages (screening 96.5%, diagnosis 96.1%, treatment 95.8%), suggesting the LLM-generated content is not systematically biased toward any single stage. The treatment stage has the lowest pre-revision *answer correctness* score, 86.1%, because biomarker and prognostic claims are easier for the LLM to over-generate than morphological descriptions; tightening the treatment-stage prompts (constraining biomarker mentions to those present in the source pathology report) raised this dimension to 96.4% post-revision. The per-stage record counts of the final pipeline export are summarised in Table 6 (1,068,187 screening, 680,078 diagnosis, 101,585 treatment-stage pairs).

Limitations of this audit. Three specialists, while sufficient for κ to be meaningful, is below the panel sizes used in formal radiology guideline development; we therefore treat this audit as a quality-assurance pass rather than a clinical gold standard. The full auditing log – per-record flag, dimension scored, and (where applicable) the resulting prompt revision – is released alongside the dataset to allow downstream re-auditing.

E.4 Baseline Model Configuration

All baseline models are evaluated zero-shot with the same instruction-formatted prompts used for BreastGPT. For proprietary models (GPT-5.4, Claude-opus-4-6, Claude-sonnet-4-6, Gemini-3.1-Flash, Gemini-3.1-Pro), we use the official APIs with temperature 0 and max_tokens 512. For open-source models we use Hugging Face Transformers with greedy decoding. Input images are preprocessed according to each model's recommended processor; for multi-image or multi-sequence cases all models receive the same ordered visual inputs and the same textual task context.

F Additional Results and Ablations

F.1 Per-Modality Performance Breakdown

Table 9 provides a per-modality performance breakdown for BreastGPT and the strongest competing baselines. This view complements the stage-level tables in the main paper by isolating whether improvements are driven by a single easy modality or are distributed across heterogeneous imaging sources.

Model	BUS	CT	Mammo	MRI	Histo	Avg
BreastGPT (cluster)	**86.81**	77.21	**75.00**	**82.86**	**71.38**	**78.65**
GPT-5.4	64.89	**78.55**	68.51	41.43	32.28	57.13
Gemini-3.1-Pro	68.09	73.33	50.21	47.14	46.53	57.06
Lingshu	58.94	**78.55**	39.89	54.29	51.52	56.64

Table 9: Per-modality performance breakdown on BreastStage-Bench closed-ended VQA.

F.2 Extended Baseline Comparisons

Tables 10 and 11 report the same VQA and Caption/Report metrics as the main paper's Sections 5 and 5.3, but extended to eight additional baseline models (Grok-4.1-Fast, Gemma-4, GLM-4.6V-Flash, LLaVA-OneVision-1.5, HealthGPT, Hulu-Med, MedDr, RadFM) covering both proprietary and recently released open / medical VLMs. None of these baselines beat BreastGPT on any cell, so for compactness the main paper limits its comparison to the original cohort; this section confirms that the conclusion holds against a broader set.

F.3 Open-ended VQA Raw Metric Breakdown

Table 12 reports the three constituent metrics (BERTScore F1, BLEU, ROUGE-1) underlying the weighted score reported in Section 5. Closed-ended VQA is omitted because its only metric is accuracy.

F.4 3D Grounding IoU on CT and MRI

None of the 21 baseline models in our evaluation cohort can localise lesions in 3D volumetric modalities (CT, MRI). Their bbox outputs are uniformly 2-D coordinate quadruples, which are dimensionally incompatible with the 6-element 3-D ground-truth boxes; consequently, on every GT-positive 3-D sample they fail to overlap the ground truth at all. The grounding-IoU numbers a few baselines report on these modalities are entirely a by-product of true-negative credit on the no-abnormality samples, not of any spatial localisation. We therefore report 3-D grounding IoU only for BreastGPT and its controlled SFT counterpart in Table 13.

Model	#P	CT IoU	MRI IoU
Qwen3-VL (SFT)	**8B**	3.89	11.98
BreastGPT (cluster)	**8B**	**5.12**	**33.49**

Table 13: Grounding IoU (%) on the 3D modalities (CT, MRI), restricted to the two models that emit volumetric (6-D) bounding boxes. IoU credits true negatives as 1.0 and penalises false positives as 0. All 21 other baselines produce 2-D outputs that are dimensionally incompatible with the 3-D ground truth and are omitted here.

F.5 BreastGPT Grounding Recognition Breakdown

Table 14 decomposes the IoU column of Section 5.3 for BreastGPT (cluster) into its true-negative and false-positive components: TN counts items with no GT bbox where the model also abstains (correct "no abnormality"), FP counts items where the model hallucinates a bbox on a normal case, and *Pred-with-GT* counts items where BreastGPT attempted a bbox on a GT-positive case.

Modality	IoU (%)	TN	FP	Pred-with-GT	N
BUS	79.59	0	0	1000	1000
CT	5.12	33	70	333	510
Mammo	23.14	0	2	998	1000
MRI	33.49	108	60	202	400

Table 14: BreastGPT (cluster) grounding-IoU decomposition. IoU credits true negatives (TN) as 1.0 and penalises false positives (FP) as 0. *Pred-with-GT* counts items where BreastGPT attempted a bbox on a GT-positive case; N is the total number of items.

F.6 Inference Efficiency Analysis

We profile the actual cost of running BreastGPT on the histopathology WSI subset of BreastStage-Bench at different token budgets k, including a

		Closed-ended VQA (Accuracy, %)										Open-ended VQA (normalized Score, %)							
		Screening				Diagnosis			Treatment			Screening			Diagnosis		Treatment		
Model	#P	BUS	CT	Mam	MRI	BUS	Mam	MRI	MRI	His	Avg	BUS	CT	MRI	BUS	MRI	MRI	His	Avg
Proprietary Models																			
Grok-4.1-Fast	–	57.23	76.61	30.74	40.95	**49.47**	**22.82**	44.23	31.75	36.42	43.36	44.76	43.43	43.95	43.06	43.94	41.74	40.50	43.05
Open-Source Models																			
Gemma-4	8B	**62.98**	**78.18**	44.57	38.57	47.87	18.17	48.43	39.68	40.68	46.57	43.29	44.20	43.46	42.47	44.49	42.40	40.12	42.92
GLM-4.6V-Flash	9B	58.94	73.09	37.34	41.90	45.48	12.16	48.60	**53.97**	49.70	46.80	42.69	43.70	42.78	42.05	43.54	42.47	40.28	42.50
LLaVA-OneVision-1.5	8B	58.09	68.85	30.32	40.95	41.76	12.31	**50.17**	**53.97**	**50.18**	45.18	–	56.06	59.84	–	59.19	58.12	45.25	55.69
Medical-Specific Models																			
HealthGPT	32B	47.02	24.48	41.38	45.71	34.04	21.17	43.01	42.86	20.22	35.54	54.66	50.17	53.08	51.92	59.38	55.27	46.41	52.98
Hulu-Med	7B	**62.98**	77.09	**48.19**	31.90	49.20	15.62	41.43	50.79	**50.18**	**47.49**	58.98	60.84	**75.46**	55.95	72.28	59.48	50.18	61.88
MedDr	40B	48.51	47.39	25.21	**52.38**	35.37	22.52	49.83	43.65	39.46	40.48	**61.46**	**64.52**	67.68	**62.83**	**73.03**	64.86	**50.46**	**63.55**
RadFM	14B	26.81	39.39	16.38	36.67	23.94	5.86	17.13	11.90	4.63	20.30	50.51	61.07	67.63	50.80	69.15	**66.76**	50.44	59.48

Table 10: VQA performance (%) on BreastStage-Bench for eight additional baselines (extends Section 5). Same column structure as the main table; refer to it for column legend.

		Caption							Report
		BUS		CT	Mammo		Histo	MRI	MRI
Model	#P	IoU	Wtd	Wtd	IoU	Wtd	Wtd	Wtd	Wtd
Proprietary Models									
Grok-4.1-Fast	–	13.40	47.30	45.12	2.52	48.09	45.15	**49.02**	50.47
Open-Source Models									
Gemma-4	8B	**37.82**	46.34	44.45	**8.01**	46.33	44.04	47.87	49.16
GLM-4.6V-Flash	9B	28.03	43.01	43.43	6.01	43.15	43.53	44.26	48.21
LLaVA-OneVision-1.5	8B	14.99	47.68	49.04	4.53	48.41	46.91	48.62	**50.62**
Medical-Specific Models									
HealthGPT	32B	22.56	43.11	40.05	7.69	44.40	46.63	46.93	48.18
Hulu-Med	7B	5.75	**48.67**	48.40	0.41	**48.74**	**47.25**	48.49	50.33
MedDr	40B	0.17	47.68	**50.25**	0.38	48.34	44.73	47.04	46.60
RadFM	14B	-	41.90	38.13	-	43.42	41.23	46.53	43.27

Table 11: Caption and Report Generation Performance (%) for eight additional baselines (extends Section 5.3). Same column structure as the main table.

		Screening			Diagnosis		Treatment	
Model	#P	BUS	CT	MRI	BUS	MRI	MRI	His
Proprietary Models								
GPT-5.4	–	88.88/1.62/34.33	86.50/1.65/22.73	91.82/11.65/44.10	87.25/1.45/25.71	90.82/14.09/41.05	90.57/10.10/41.04	83.67/0.57/11.02
Claude-opus-4-6	–	82.55/0.45/7.73	81.86/0.66/6.70	82.75/1.25/8.24	82.05/0.30/6.81	82.95/1.87/9.90	81.76/0.73/7.40	79.96/0.17/3.27
Claude-sonnet-4-6	–	82.54/0.46/8.97	81.71/0.88/8.16	82.71/1.26/10.56	81.85/0.34/7.94	82.83/1.82/12.46	80.90/0.51/8.15	79.90/0.14/4.32
Gemini-3.1-Flash	–	86.50/1.51/20.51	86.04/2.68/18.08	85.60/2.51/15.81	84.66/0.68/13.19	85.36/3.26/16.15	83.07/1.12/9.34	82.27/0.34/6.37
Grok-4.1-Fast	–	83.70/0.56/11.10	82.32/0.82/8.25	82.83/0.90/9.24	82.21/0.30/7.51	82.33/1.29/9.83	80.84/0.34/4.93	79.63/0.13/2.60
Open-Source Models								
Qwen2.5-VL	3B	87.99/2.12/19.97	86.85/2.98/16.89	87.59/3.41/18.86	86.69/1.31/14.75	87.48/5.50/21.70	85.49/2.08/15.08	84.46/0.49/9.80
Qwen2.5-VL	7B	86.11/1.13/11.45	85.82/2.08/13.72	87.36/2.70/20.45	85.06/0.48/9.36	86.98/4.61/20.11	83.92/1.32/9.62	83.34/0.32/8.94
Qwen3-VL	4B	84.69/0.98/10.42	89.65/5.08/37.59	86.29/2.10/16.16	84.23/0.91/10.36	87.78/4.95/25.31	85.61/1.86/22.86	80.65/0.22/3.52
Qwen3-VL	8B	84.40/1.00/9.98	86.40/2.64/18.68	83.89/1.73/9.90	83.76/0.74/9.11	84.43/2.78/13.30	82.52/1.14/10.41	80.51/0.23/3.60
MiMo-VL	7B	82.27/0.52/4.90	88.14/4.29/43.35	89.97/19.76/55.10	81.32/0.37/3.75	90.36/22.57/59.33	89.91/16.07/56.42	78.85/0.14/1.59
InternVL3.5	8B	87.86/4.39/22.84	89.02/14.39/39.77	90.19/7.99/37.85	86.56/2.69/19.37	90.65/12.68/41.59	89.44/7.36/37.54	84.86/0.50/15.80
Gemma-4	8B	82.79/0.52/7.05	83.20/1.12/9.26	82.71/1.05/7.39	81.86/0.32/5.83	83.17/1.91/9.71	81.59/0.68/5.75	79.20/0.05/2.04
GLM-4.6V-Flash	9B	82.16/0.71/5.73	83.03/1.35/7.38	82.38/1.16/5.20	81.52/0.57/4.58	82.90/1.69/6.65	82.17/0.89/4.65	79.42/0.19/2.10
LLaVA-OneVision-1.5	8B	–	90.14/6.57/37.39	91.82/11.65/44.09	–	90.82/14.09/41.03	90.57/10.10/41.24	84.88/0.39/10.84
Medical-Specific Models								
Lingshu	7B	89.47/3.66/27.78	87.79/4.01/19.76	89.29/5.04/30.02	87.59/1.80/18.81	88.65/8.22/29.52	86.70/3.17/21.84	83.57/0.30/7.32
HuatuoGPT-V	7B	88.99/3.60/25.86	88.43/5.17/23.85	90.20/6.84/35.48	87.66/2.25/19.23	89.36/10.04/31.58	88.17/5.16/30.22	85.07/0.88/12.08
HealthGPT	32B	89.69/5.73/33.52	86.17/7.50/20.84	87.02/10.10/28.18	88.90/3.75/26.11	89.91/16.43/41.27	87.84/10.10/35.30	84.83/0.71/15.28
Hulu-Med	7B	91.14/9.89/43.73	90.99/14.81/46.56	95.29/40.64/70.61	90.02/6.34/37.42	93.60/37.93/64.00	90.64/10.31/46.35	87.35/3.33/22.70
MedDr	40B	91.22/20.02/43.39	90.46/31.04/46.13	92.69/26.61/58.74	91.85/19.38/48.23	93.70/40.07/64.63	91.82/21.43/54.38	87.12/3.25/24.35
RadFM	14B	87.50/3.86/23.19	89.66/24.94/40.04	93.08/25.80/58.57	86.53/9.71/20.44	92.94/30.76/59.96	91.86/26.74/56.57	87.00/3.21/24.55
Qwen3-VL (SFT)	**8B**	**98.85/87.18/92.83**	**99.04/89.22/94.24**	**99.11/85.25/94.82**	**98.26/81.90/88.38**	**99.03/87.90/95.17**	**97.65/72.02/86.10**	**90.73/11.46/38.93**
BreastGPT (cluster)	**8B**	**99.15/90.67/94.91**	**99.07/88.56/94.45**	**99.27/87.53/95.87**	**98.63/85.08/90.60**	**99.12/88.94/95.68**	**98.09/75.47/88.08**	**92.69/18.88/50.96**

Table 12: Open-ended VQA raw metrics per cell. Each entry is BERTScore F1 / BLEU / ROUGE-1 (%). BreastGPT in cyan.

no-limit baseline that disables the selector and feeds every patch token of the slide into the LLM. Protocol: bf16, batch size 1, 5 warm-up + 30 timed forward passes, on a single GPU. La-

tency is measured with `torch.cuda.Event` and `torch.cuda.synchronize()` so that timing reflects the actual GPU execution rather than asynchronous launches. Selector latency is the wall-clock time inside the greedy coverage routine; prefill latency is the first-token forward through the LLM after selection; total latency is selector + prefill.

Three observations follow from these numbers; the per-stage latency breakdown is plotted in main-paper Fig. 5. *(i) Prefill and decode scale roughly linearly with k.* From k=1 to k=512 the LLM input length grows from 145 to 656 tokens, and prefill grows in step from 57.9 to 605.7 ms; decode time per output token grows from 42.6 to 150.7 ms, since the KV-cache that each new token must attend to also lengthens with k. The chosen k=128 costs 191 ms of prefill, 3.2× less than k=512, at >99% of its task quality. *(ii) Without selection, the LLM forward dominates the entire pipeline.* The no-limit baseline raises the LLM input to 5,144 tokens; prefill jumps to 5.5 s and per-token decode to 1.1 s, for a total of 6.5 s per question—about 33× slower than k=128 on the same slide. The selector itself only takes 1.0 s in this regime, so the cost saved is overwhelmingly LLM-side, exactly as expected for a method that compresses the LLM input. *(iii) Memory grows slowly under selection and jumps only at no limit.* Peak GPU memory rises from 16.96 GB at k=1 to 16.98 GB at k=512 and then to 17.95 GB without selection, a ∼1 GB jump that scales with the LLM input length. BreastGPT therefore runs on a single 24 GB GPU at any selected budget; the no-limit configuration would exceed that envelope on larger slides.

G Detailed Ablation Results

G.1 Numerical Results for WSI Branch Ablation

Table 16 reports an internal design ablation of the GigaPixel branch, complementary to the Qwen3-VL (SFT, 32-patch WSI) baseline reported in the main paper. This ablation isolates the contribution of each subcomponent inside the GigaPixel branch itself: average pooling loses spatial heterogeneity, truncation ignores sparse diagnostic regions, V-V-only selection may miss query-relevant evidence, and T-V-only selection may overfocus on a small salient region.

The results demonstrate that each component contributes incrementally. LongNet contextualization improves closed-ended accuracy by 9.5 points over average pooling, showing that slide-level context is essential for WSI reasoning. Adding V-V coverage gains another 6.6 points, indicating that global representativeness matters beyond sequential context aggregation. Incorporating T-V coverage yields the final 3.8-point improvement, confirming that clinical-query alignment is necessary once the selector is asked to compress the slide into a very small token budget.

G.2 Numerical Results for Visual Token Budget Sweep

Tables 17 and 18 provide the complete numerical breakdown of the visual token budget sweep summarized in main-paper Fig. 4. We sweep $k \in \{1, 8, 16, 32, 64, 128, 256, 512\}$ and report performance on every modality and every task type in BreastStage-Bench: VQA (closed- and open-ended) across screening, diagnosis, and treatment stages, as well as caption and report generation. This is intentionally exhaustive so that the operating-point choice (k=128) used throughout the main paper can be audited per-modality, not only on histopathology.

Two patterns hold consistently across modalities. *(i) Performance saturates at k=128.* On VQA closed-ended screening, BUS reaches its plateau at k=64 (86.81%) and stays flat through k=512; CT, Mammo, and MRI similarly fluctuate within ±1 point past k=64. On caption generation, the weighted score on BUS rises from 71.01 at k=1 to 79.32 at k=128 and gains only 0.06 points by k=512; Mammo gains only 0.21 points from 128 to 512. On open-ended VQA, all six (modality, stage) cells are within 0.5 points of their maximum by k=128. *(ii) The ground caption task needs a minimum number of tokens.* BUS ground caption IoU jumps from 1.50 at k=1 to 79.59 at k=128 before plateauing, and Mammo ground caption IoU from 0.85 to 23.14. This indicates that a single coverage-selected token is insufficient to anchor a lesion bounding region, but k=128 already recovers most of the IoU achievable at k=512. The 128-token budget therefore gives BreastGPT the same downstream inference cost on radiology and pathology while staying on the saturation plateau of every metric we measure.

k	LLM in. tok.	Prefill (ms)	Selector (ms)	Decode (ms/tok)	Total (ms)	Peak Mem (GB)
1	145	57.9	0.4	42.6	58.2	16.96
8	152	67.3	0.9	45.9	68.2	16.96
16	160	78.1	1.5	50.4	79.6	16.96
32	176	89.4	2.6	48.0	91.9	16.96
64	208	123.1	4.8	54.5	127.9	16.96
128	272	191.4	9.3	68.2	200.6	16.97
256	400	328.4	18.2	95.5	346.6	16.97
512	656	605.7	36.1	150.7	641.8	16.98
no limit	5,144	5,512.6	—	1,129.5	6,545.7	17.95

Table 15: Inference efficiency of BreastGPT (cluster) on a representative histopathology WSI (5,987 patches before selection). "LLM in. tok." is the total LLM input length after selection (visual k tokens plus prompt). The "no limit" row sets k above the patch count, effectively disabling the selector and feeding every patch token to the LLM.

Configuration	Closed Acc (%)	Caption BERT	Caption BLEU	Caption R-1
CONCH + LongNet + V-V only	69.4	0.86	0.033	0.248
CONCH + LongNet + T-V + V-V (full)	**71.4**	**0.88**	**0.038**	**0.267**

Table 16: Numerical results for WSI branch ablation study on histopathology tasks.

Token	Screen							Diagnosis					Treatment			
	closed				open			closed			open		closed		open	
	BUS	CT	Mammo	MRI	BUS	CT	MRI	BUS	Mammo	MRI	BUS	MRI	MRI	Histo.	MRI	Histo.
1	60.21	78.79	43.41	72.38	91.84	95.66	94.80	51.06	48.80	65.39	87.82	94.63	40.48	66.63	89.63	56.82
8	75.32	78.43	60.43	80.95	93.76	94.89	95.47	69.94	64.41	70.28	90.53	94.99	46.83	71.25	89.22	63.78
16	84.04	77.70	72.88	81.43	95.41	94.83	95.83	75.26	67.42	78.32	92.02	95.26	59.52	71.01	89.67	66.12
32	86.17	77.70	74.05	85.24	95.80	95.03	95.58	76.86	68.92	81.29	93.08	95.71	62.70	71.01	89.96	66.53
64	86.81	76.73	74.69	83.81	95.96	95.48	95.58	77.66	68.77	81.12	92.94	95.72	61.11	71.13	89.95	66.08
128	86.81	77.21	75.00	82.86	95.97	95.29	95.48	77.13	68.32	81.12	93.24	95.72	61.11	71.38	89.93	63.80
256	86.81	76.73	74.05	82.38	95.97	95.51	95.75	77.13	68.02	81.29	93.24	95.61	59.52	71.01	89.97	66.05
512	86.81	77.58	74.47	81.90	95.96	95.70	95.65	77.13	68.32	82.69	93.23	95.67	57.14	71.50	89.89	64.58

Table 17: Token budget sweep on the VQA task across screening, diagnosis, and treatment stages, under closed-ended and open-ended settings. "Histo." denotes histopathology.

Token	Caption															Report									
	BUS					CT					Mammo					Histopathology					MRI				
	BERT-F1	BLEU	R-1	Wtd	IoU	BERT-F1	BLEU	R-1	Wtd	IoU	BERT-F1	BLEU	R-1	Wtd	IoU	BERT-F1	BLEU	R-1	Wtd		BERT-F1	BLEU	R-1	Wtd	IoU
1	92.77	36.61	61.89	71.01	1.50	93.90	38.01	65.46	72.83	4.90	93.77	38.35	64.75	72.66	0.85	83.86	18.23	32.70	54.67		89.85	26.01	54.91	65.15	24.52
8	94.08	45.79	69.41	75.84	20.00	93.70	37.27	64.40	72.28	4.91	94.20	44.52	67.53	75.11	2.13	88.34	40.47	53.35	67.63		90.09	27.27	55.43	65.71	24.72
16	94.63	50.24	72.44	77.99	40.88	93.77	38.18	64.69	72.61	4.93	94.67	49.48	69.88	77.18	5.00	88.04	39.57	50.99	66.66		90.37	28.82	56.52	66.51	25.41
32	94.90	52.21	73.85	78.97	63.04	93.88	38.92	65.34	73.01	4.93	94.72	50.04	70.03	77.38	12.52	87.82	37.51	49.98	65.79		90.50	30.07	57.47	67.13	26.17
64	94.94	52.71	74.23	79.21	76.76	93.89	40.89	65.89	73.65	5.03	94.73	50.35	70.05	77.46	18.37	88.15	39.50	51.81	66.91		90.58	30.90	58.11	67.53	28.55
128	94.97	52.91	74.42	79.32	79.59	93.76	39.81	65.30	73.16	5.12	94.75	50.74	70.31	77.64	23.14	88.12	39.61	51.25	66.78		90.61	31.24	58.24	67.67	33.49
256	95.00	53.48	74.55	79.51	80.33	93.73	39.41	65.24	73.03	5.03	94.84	51.57	70.91	78.04	23.06	88.00	37.67	50.14	65.96		90.59	31.16	58.09	67.60	29.64
512	94.98	53.10	74.45	79.38	80.36	93.79	39.96	65.61	73.29	5.02	94.80	51.11	70.69	77.85	23.10	87.58	34.94	47.98	64.53		90.62	31.21	58.29	67.68	29.52

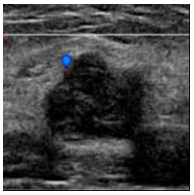
Table 18: Token budget sweep on the caption and report generation tasks. "Wtd" denotes the weighted average of BERT-F1, BLEU, and ROUGE-1. The "IoU" column reports the *Ground Caption* task (lesion-grounded captioning), which is only defined for the modalities with bounding-region annotations.

H Qualitative Analysis and Case Studies

H.1 Case Analysis

To complement the aggregate metrics in Sections 5 and 5.3, we walk through six representative cases drawn from BreastStage-Bench that span all five modalities and all four task families (closed-ended VQA, open-ended VQA, lesion-grounded captioning, and structured report generation). For each case we show the role-conditioned system prompt that BreastGPT receives at inference time (abbreviated to its lead clause), the user instruction, and the reference output that defines the gold answer for the metrics in the main paper. Together they il-

lustrate the breadth of input formats and reasoning styles a single BreastGPT checkpoint must support.



Case 1 – Screening-stage Closed VQA

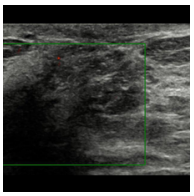
System. You are an expert screening breast radiologist. Your clinical focus is on early detection and risk triage across broad populations. You possess a high sensitivity for identifying subtle abnormalities within varied anatomical backgrounds to determine if further investigation is warranted. Now, apply your expertise to the following task: Answer the multiple-choice question based on the provided breast ultrasound image(s). Respond ONLY with the exact option chosen.

User. `<image>`
Classify the blood flow status: No flow, Mild flow, or Abundant flow?
Options: (A) No flow (B) Abundant flow (C) Mild flow
Answer:

Reference output. (C) Mild flow

Comment. The screening persona drives the model toward a categorical Doppler readout rather than a free-text discussion, and the answer must come from the schema enum so it can be scored by exact-string accuracy.

Figure 11: Case 1: closed-ended vascularity attribute on a single BUS B-mode frame from BUS-CoT.



Case 2 – Screening-stage Open-ended VQA

System. You are an expert screening breast radiologist. Your clinical focus is on early detection and risk triage across broad populations. You possess a high sensitivity for identifying subtle abnormalities within varied anatomical backgrounds to determine if further investigation is warranted. Now, apply your expertise to the following task: Answer the open-ended question based on the provided breast ultrasound image(s).

User. `<image>`
Describe the degree of vascularity observed within the lesion on Doppler ultrasound.

Reference output. Mild flow is detected within the lesion.

Comment. The same underlying clinical fact as Case 1 is rewritten into a fluent description; this pair is what allows the weighted (BERTScore F1, BLEU, ROUGE-1) score to evaluate semantic equivalence rather than exact-string match.

Figure 12: Case 2: open-ended Doppler description on a BUS Doppler frame from BUS-CoT.



Case 3 – Lesion-grounded Captioning

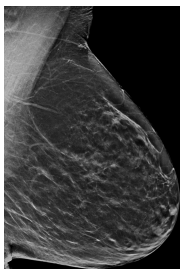
System. You are an expert breast clinician. Your clinical focus is on comprehensive breast disease assessment and patient-centered clinical management. You integrate clinical presentation, imaging findings, pathological evidence, risk factors, and treatment considerations to provide balanced, clinically grounded assessments across the full breast care pathway. Now, apply your expertise to the following task: Generate a descriptive caption for the provided breast ultrasound image(s). You MUST explicitly ground each mentioned finding with a 2D bounding box [xmin, ymin, xmax, ymax].

User. `<image>`
Locate any abnormalities in this breast ultrasound and provide a detailed description of their morphological features.

Reference output. Standard Ultrasound (B-Mode) reveals `<ref-object><bbox>`. This finding demonstrates microcalcifications of suspicious morphology and is categorized as BI-RADS 4A, prompting a recall for biopsy or further evaluation. Despite benign morphological features, the presence of suspicious calcifications elevates concern, though the overall prognosis remains favorable with routine follow-up.

Comment. The expected output interleaves a 2D bounding box with a narrative description, exercising both spatial localisation (scored by IoU in Section 5.3) and BI-RADS-level reasoning (scored by the weighted text metrics).

Figure 13: Case 3: lesion-grounded captioning on a BUS expert-cropped frame from BUS-CoT/BUS-Expert.



Case 4 – Diagnosis-stage Closed VQA (BI-RADS)

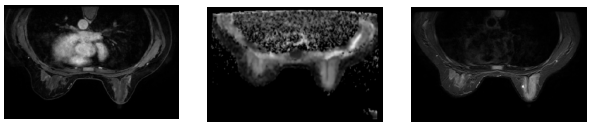
System. You are an expert diagnostic breast radiologist. Your clinical focus is on definitive lesion characterization and differential diagnosis. You rigorously analyze imaging features to separate benign from malignant entities, providing conclusive assessments that directly guide patient management. Now, apply your expertise to the following task: Answer the multiple-choice question based on the provided mammogram image(s). Respond ONLY with the exact option chosen.

User. `<image>`$\times 8$
Can you provide the Bi-Rads category for this mammography exam?
Options: (A) Bi-Rads 3 (B) Bi-Rads 6 (C) Bi-Rads 0 (D) Bi-Rads 1 (E) Bi-Rads 2 (F) Bi-Rads 4 (G) Bi-Rads 5
Answer:

Reference output. (D) Bi-Rads 1

Comment. The eight-view input forces multi-image fusion within a single forward pass; the output also illustrates BreastStage's commitment to the full BI-RADS 0–6 enum, including the assessment-incomplete category (0) and the established-malignancy category (6) that purely diagnostic benchmarks frequently omit.

Figure 14: Case 4: BI-RADS classification on an 8-view mammographic exam from EMBED (one view shown).



T1dyn ADC T2w

Case 5 – Screening-stage Report Generation

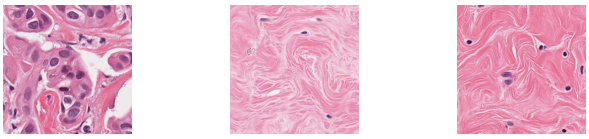
System. You are an expert breast clinician. Your clinical focus is on comprehensive breast disease assessment and patient-centered clinical management. You integrate clinical presentation, imaging findings, pathological evidence, risk factors, and treatment considerations to provide balanced, clinically grounded assessments across the full breast care pathway. Now, apply your expertise to the following task: Generate a comprehensive clinical report based on the provided multi-sequence breast MRI.

User. Breast MRI sequences: ADC: `<video>`. DWI: `<video>`. T1: `<video>`. T1dyn: `<video>`. T2w: `<video>`.
Based on the provided breast MRI images, generate a comprehensive and professional MRI report detailing all relevant findings, diagnostic impression with BI-RADS classification, and appropriate clinical recommendations.

Reference output. Breast MRI Report.
Findings: The bilateral breasts are symmetric in size and contour, with heterogeneously dense (ACR Category C) fibroglandular tissue and mild background parenchymal enhancement. No architectural distortion or abnormal parenchymal distribution is identified. In the posteroinferior inner quadrant of the left breast, there is an irregular, spiculated mass measuring approximately 13 mm × 11 mm × 16 mm. This lesion demonstrates marked early enhancement with washout kinetics and shows restricted diffusion on DWI/ADC sequences. No additional suspicious enhancing foci are seen in the left or right breast. Bilateral skin thickness is normal, nipples are symmetric without retraction, and there is no evidence of chest wall invasion. Axillary lymph nodes bilaterally appear morphologically benign without pathological enlargement. An incidental finding of a simple hepatic cyst is noted in the liver.
Impression: Irregular, markedly enhancing mass in the posteroinferior inner quadrant of the left breast with diffusion restriction, consistent with known malignancy. BI-RADS 6. Right breast shows no MRI evidence of malignancy. BI-RADS 1. Incidental simple hepatic cyst, likely benign.
Recommendation: Given the BI-RADS 6 classification, this finding reflects a known or biopsy-proven malignancy. Proceed with definitive surgical management or neoadjuvant therapy planning as per multidisciplinary consensus. No further diagnostic imaging of the breast is required for this lesion. The hepatic cyst requires no intervention or follow-up.

Comment. The report task simultaneously requires multi-sequence integration, structured section ordering (Findings → Impression → BI-RADS → Management), and quantitative descriptors that must be supported by the actual volumes; failures concentrate here in our Section 5.3 numbers.

Figure 15: Case 5: comprehensive multiparametric MRI report from FUDAN; mid-slice frames of three of the five paired sequences shown above.



patch 1 patch 9 patch 17

Case 6 – Pathology Report Generation

System. You are an expert breast clinician. Your clinical focus is on comprehensive breast disease assessment and patient-centered clinical management. You integrate clinical presentation, imaging findings, pathological evidence, risk factors, and treatment considerations to provide balanced, clinically grounded assessments across the full breast care pathway. Now, apply your expertise to the following task: Generate a concise descriptive caption for the provided breast histopathological image(s).

User. `<image>×5024`
Please generate a comprehensive and professional pathology report based on the provided slide from Breast, including key elements such as: (1) the histopathological diagnosis, (2) tumor morphology and growth patterns, (3) extent of invasion and involved structures, (4) lymphovascular or perineural invasion, (5) status of surgical margins, (6) TNM staging if possible, and (7) any additional pathological findings including non-neoplastic tissue alterations.

Reference output. The pathology findings indicate a diagnosis of infiltrating ductal carcinoma located in the upper outer quadrant of the left breast. The tumor exhibits a moderately differentiated histologic grade, suggesting an intermediate level of aggression and atypia. There is grossly evident lesion upon inspection. Importantly, the surgical margins are uninvolved, indicating no residual tumor at the edges of the excised tissue. Lymph node analysis reveals that all ten axillary lymph nodes examined are negative for metastatic involvement, suggesting localized disease without regional spread. The Nottingham score, which typically helps in grading the tumor, could not be determined in this case. Overall, the pathological evaluation supports a moderately aggressive tumor confined to the primary site with no lymphatic dissemination.

Comment. On gigapixel inputs the GigaPixel branch must compress thousands of patch tokens down to k=128 before the LLM ever sees them; the section structure of the reference output (diagnosis → grade → margins → nodes → TNM) is exactly what the coverage selector is asked to preserve.

Figure 16: Case 6: pathology report on a TCGA-BRCA whole-slide image at 20× magnification; 3 of the 5024 sampled patches shown above.

What the cases collectively show. The hard cases for BreastGPT cluster around three axes: (1) fine-grained categorical boundaries when the schema admits adjacent values (BI-RADS 3 vs. 4A, Nottingham grade II vs. III); (2) quantitative predictions that should be supported by an explicit measurement but are inferred from the visual signal alone (ADC values, Ki-67 index); and (3) molecular-biomarker inference (ER/PR/HER2 status) from H&E morphology alone, without paired immunohistochemistry. These observations motivate future work on uncertainty-aware reporting, calibrated abstention on adjacent-category questions, and explicit integration of non-image clinical evidence at inference time.