

Historical Report Guided Bi-modal Concurrent Learning for Pathology Report Generation

Ling Zhang¹, Boxiang Yun¹, Qingli Li¹, and Yan Wang¹(✉)

Shanghai Key Laboratory of Multidimensional Information Processing,
East China Normal University, Shanghai, China
51265904017@stu.ecnu.edu.cn, 52265904012@stu.ecnu.edu.cn,
qlii@cs.ecnu.edu.cn, ywang@cee.ecnu.edu.cn

Abstract. Automated pathology report generation from Whole Slide Images (WSIs) faces two key challenges: (1) lack of semantic content in visual features and (2) inherent information redundancy in WSIs. To address these issues, we propose a novel Historical Report Guided **Bi**-modal Concurrent Learning Framework for Pathology Report **Generation** (Bi-Gen) emulating pathologists’ diagnostic reasoning, consisting of: (1) A knowledge retrieval mechanism to provide rich semantic content, which retrieves WSI-relevant knowledge from pre-built medical knowledge bank by matching high-attention patches and (2) A bi-modal concurrent learning strategy instantiated via a learnable visual token and a learnable textual token to dynamically extract key visual features and retrieved knowledge, where weight-shared layers enable cross-modal alignment between visual features and knowledge features. Our multi-modal decoder integrates both modalities for comprehensive diagnostic reports generation. Experiments on the PathText (BRCA) dataset demonstrate our framework’s superiority, achieving state-of-the-art performance with 7.4% relative improvement in NLP metrics and 19.1% enhancement in classification metrics for Her-2 prediction versus existing methods. Ablation studies validate the necessity of our proposed modules, highlighting our method’s ability to provide WSI-relevant rich semantic content and suppress information redundancy in WSIs. Code is publicly available at <https://github.com/DeepMed-Lab-ECNU/BiGen>.

Keywords: Whole Slide Image · Image Caption · Pathology Report Generation.

1 Introduction

The automated analysis of Whole Slide Images (WSIs) in digital pathology has emerged as a significant research direction in the field of medical artificial intelligence recently [7,17,26]. Among these, the task of pathology report generation based on WSIs has garnered considerable attention due to its immense potential in assisting diagnosis and reducing the workload of physicians [3,10,21,22].

In recent years, there have been a significant number of studies on image captioning [15,25,27,29]. In the medical field, numerous studies have focused

on generating radiology reports from radiological images [5,6,9]. There are also studies [12,16,23] that have focused on obtaining descriptions of small pathological regions of interest. Compared with these report generation tasks, generating reports from WSIs is more challenging and less studied. The challenge relies mainly on how to effectively extract **semantically diagnostic features** from WSIs with **ultra-high resolution** [13,28]. Existing approaches such as MI-Gen [3] directly apply a Transformer model to input image tokens, allowing each token to interact with other tokens. HistGen [10] employs local-global hierarchical visual encoding to input image tokens, which still relies on a self-attention mechanism to explore relationships among multiple tokens. Though global visual features can be learned via self-attention, these methods only learn visual features from WSI itself, which inevitably (1) **lack of semantic content** and (2) **introduce considerable redundancy**.

To address these challenges, we propose to leverage historical diagnostic reports to guide the learning of semantically key pathological features for generating current WSI’s report. Concretely, we let the visual features be learned along with the relevant textual features, whose knowledge is retrieved from a pre-built knowledge bank. Our method contains two key components: (1) We propose to retrieve WSI-relevant report candidates from a pre-built knowledge bank to provide **semantically rich** diagnostic report information. (2) We propose a bi-modal concurrent learning strategy. It is instantiated via a learnable visual token and a learnable textual token which extract **key pathological features** and key knowledge features, respectively. The strategy also enables the learning of visual token to be guided by the semantically rich textual token. The knowledge retrieval emulates the process that pathologists recall historical diagnosis records from similar cases to aid current decision-making. The visual and textual token represent a more global perspective to mirror the way pathologists view WSIs and recall historical diagnosis records—they don’t necessarily examine each patch or each record but rather focus on key information to make a comprehensive diagnosis. The main contributions of this paper can be summarized as follows:

- To the best of our knowledge, we first introduce an explicit knowledge retrieval mechanism into the pathology report generation task to enhance semantically rich diagnostic report information, establishing an interpretable cross-modal knowledge transfer paradigm.
- We propose a bi-modal concurrent learning strategy to dynamically extract key features in WSIs and retrieved knowledge, effectively suppressing information redundancy while enhancing pathological semantic learning.
- Experiments on the PathText (BRCA) dataset demonstrate that our method significantly outperforms existing benchmarks in report generation quality (BLEU-4=0.135, ROUGE-L=0.293) and Her-2 prediction metrics (F1=0.730).

2.1 Problem Formulation

The task of pathology report generation is formulated as a sequence-to-sequence problem, where the input consists of multiple image tokens extracted from a high-dimensional pathology image, and the output is a natural language report. Let $\{\mathbf{w}_i\}_{i=1}^M$ denote the set of non-overlapping patches cropped from the WSI \mathbf{W} , where M is the number of patches. The objective is to generate the target pathology report $\mathbf{Y} = \{\mathbf{y}_n\}_{n=1}^N$, where N is the length of the report for \mathbf{W} .

A pretrained visual extractor is used to extract feature embeddings from $\{\mathbf{w}_i\}_{i=1}^M$, followed by a linear projection layer, resulting in a feature matrix $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^M$, where $\mathbf{x}_i \in \mathbb{R}^d$ and d is the embedding dimension. The training objective is to maximize the likelihood of generating the correct report given the patch embeddings \mathbf{X} . The total negative log-likelihood (NLL) loss is computed as: $\mathcal{L} = -\sum_{n=1}^N \log P(\mathbf{y}_n | \{\mathbf{y}_i\}_{i < n}, \mathbf{X})$, where $\{\mathbf{y}_i\}_{i < n}$ represents the previously generated tokens in the report sequence.

2.2 Visual Token Cross-attention

In the visual branch, we propose a visual token cross-attention module to extract a compact and informative visual representation of the WSI. Instead of performing self-attention over all patch features, we utilize a cross-attention mechanism, where a single learnable token attends to all patch features. Specifically, we initialize a visual token $\mathbf{V}_0 \in \mathbb{R}^{1 \times d}$ as the query. The patch features $\mathbf{X} \in \mathbb{R}^{M \times d}$ act as the key and the value, and the cross-attention mechanism is applied iteratively over L layers, yielding the following update rule:

$$\mathbf{V}_l = \text{CrossAttn}(\mathbf{V}_{l-1}, \mathbf{X}; \boldsymbol{\Theta}_l), l \in \{1, 2, \dots, L\}, \quad (1)$$

where $\text{CrossAttn}(\cdot)$ denotes the cross-attention layer, and $\boldsymbol{\Theta}_l$ means the l -th layer's learnable parameters.

2.3 Knowledge Branch

Pathologists frequently recall historical diagnosis records to support diagnostic decisions. Inspired by this practice, we incorporate a knowledge retrieval mechanism to enhance the pathology report generation. PLIP [11] is a visual-language foundation model trained on over 200,000 paired pathology images and text data, allowing pathologists to search for similar cases through image or text queries, which is utilized to complete knowledge retrieval in this module.

I. Knowledge Bank Construction: Due to the text token length limitation, a long report can't be fed into the text encoder of PLIP [11]. So we construct a knowledge bank from sentence-level pathology reports to provide high-quality semantic knowledge, as illustrated in Fig. 1 (b). Specifically, the original reports in the training set (*e.g.*, PathText (BRCA)) are split into individual sentences and encoded using the text encoder of PLIP [11]. The resulting sentence embeddings are stored as the knowledge bank, denoted as $\mathbf{S} = \{\mathbf{s}_i\}_{i=1}^T$, where $\mathbf{s}_i \in \mathbb{R}^d$, and T is the number of sentence embeddings in the knowledge bank.

II. Knowledge Retrieval: As illustrated in Fig. 1 (c), to retrieve knowledge using image embeddings aligned with sentence embeddings in the knowledge bank, we feed all patches $\{\mathbf{w}_i\}_{i=1}^M$ into the image encoder of PLIP [11] and obtain patch embeddings $\mathbf{X}^{\text{PLIP}} = \{\mathbf{x}_i^{\text{PLIP}}\}_{i=1}^M$, where $\mathbf{x}_i^{\text{PLIP}} \in \mathbb{R}^d$ and d is the embedding dimension. Considering that many patches may be irrelevant to diagnosis, retrieving knowledge based on all patches could introduce redundant and noisy information. Inspired by [2], which optimizes computational efficiency by learning adaptive attention patterns in early layers and removing unimportant visual tokens in subsequent layers, we use the attention scores from the first layer of the visual token cross-attention module (Sec. 2.2) to identify key patches for knowledge retrieval. Specifically, the top- k patches with the highest attention scores are selected, where k is the selecting ratio, and their corresponding image embeddings in \mathbf{X}^{PLIP} are collected as $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^{M \times k}$, where $\mathbf{p}_i \in \mathbb{R}^d$.

Although top- k patches are selected, similarity measure between selected patch embeddings and sentence embeddings still leads to massive calculation due to thousands of patch embeddings and thousands of sentence embeddings stored in the knowledge bank. Considering the spatial neighborhood patches with same tissue types on a WSI [14], we flatten the token sequence \mathbf{P} and uniformly partition it into $M \times k/m$ tissue regions with m as the region size, ensuring that sparse patch embeddings are compressed. Then the average feature for each tissue region is computed as the region feature and updated tissue region embeddings are obtained: $\bar{\mathbf{P}} = \{\bar{\mathbf{P}}_1, \bar{\mathbf{P}}_2, \dots, \bar{\mathbf{P}}_{M \times k/m}\}$.

Next, the cosine similarity between each region's feature in $\bar{\mathbf{P}}$ and all stored knowledge embeddings in the knowledge bank \mathbf{S} are computed. For the i -th region, we select the top- v knowledge features based on the similarity and obtain $\mathbf{r}_i = \{\mathbf{r}_i^j\}_{j=1}^v$, where $\mathbf{r}_i^j \in \mathbb{R}^d$, and the average knowledge feature in \mathbf{r}_i is computed as the retrieved knowledge feature $\bar{\mathbf{r}}_i$, where $\bar{\mathbf{r}}_i \in \mathbb{R}^d$. All retrieved knowledge features are obtained: $\mathbf{R} = \{\bar{\mathbf{r}}_i\}_{i=1}^{M \times k/m}$, followed by a linear projection layer.

III. Textual Token Cross-attention: To focus on key information from redundant retrieved knowledge, we introduce a learnable textual token $\mathbf{T}_0 \in \mathbb{R}^{1 \times d}$. Similar to the visual token cross-attention module, we apply a cross-attention mechanism iteratively to refine the textual global token token:

$$\mathbf{T}_l = \text{CrossAttn}(\mathbf{T}_{l-1}, \mathbf{R}; \Theta_l), l \in \{1, 2, \dots, L-1\}. \quad (2)$$

To promote cross-modal alignment, the cross-attention layers of visual and knowledge branches share weights, obtaining rich semantic content. Since the knowledge retrieval relies on the attention scores from the first visual token cross-attention, the textual token cross-attention are deployed for $L-1$ layers.

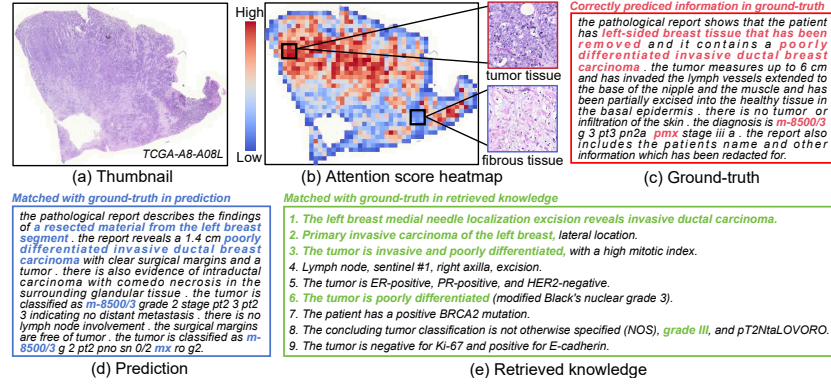
2.4 Multi-modal Decoder

After refinement for both the visual and textual tokens, we concatenate the two tokens to form a joint representation: $\mathbf{F} = \text{Concat}(\mathbf{V}_L, \mathbf{T}_{L-1})$.

The report sequence is generated auto-regressively, implemented by Masked Multi-head Self-attention. Specifically, in the l -th decoder layer, the previously

Table 1. Results of pathology report generation on PathText (BRCA). BLEU-n: the BLEU score computed based on n-grams. **Bold**: the highest score. Underline: 2nd score.

| Model | NLP metrics | | | | | | | Classification metrics | | |
|--------------------------|--------------|--------------|--------------|--------------|--------------|--------------|---------------------|------------------------|--------------|--------------|
| | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 | METEOR | ROUGE | Fact _{ent} | Precision | Recall | F1 |
| CNN-RNN [25] | 0.371 | 0.185 | 0.089 | 0.043 | 0.143 | 0.239 | 0.478 | 0.429 | 0.750 | 0.546 |
| att-LSTM [27] | 0.371 | 0.191 | 0.094 | 0.048 | 0.142 | 0.238 | 0.460 | <u>0.600</u> | 0.636 | <u>0.618</u> |
| vanilla Transformer [24] | 0.389 | 0.246 | 0.157 | 0.103 | 0.158 | 0.257 | 0.487 | 0.457 | <u>0.800</u> | 0.582 |
| R2Gen [6] | 0.378 | 0.243 | 0.160 | 0.107 | <u>0.179</u> | <u>0.279</u> | 0.505 | 0.486 | 0.630 | 0.548 |
| R2GenCMN [5] | 0.396 | 0.254 | 0.164 | 0.110 | 0.163 | <u>0.279</u> | 0.470 | 0.429 | 0.682 | 0.526 |
| MI-Gen [3] | 0.416 | 0.267 | 0.174 | 0.115 | 0.165 | 0.270 | 0.490 | 0.257 | 0.529 | 0.346 |
| HistGen [10] | <u>0.422</u> | <u>0.272</u> | <u>0.177</u> | <u>0.118</u> | 0.169 | 0.277 | <u>0.493</u> | 0.514 | 0.643 | 0.571 |
| ours | 0.450 | 0.296 | 0.196 | 0.135 | 0.180 | 0.293 | 0.521 | 0.657 | 0.821 | 0.730 |

**Fig. 2.** Visualization of the (a) thumbnail, (b) attention score heatmap, (c) report ground-truth, (d) prediction and (e) retrieved knowledge of sample "TCGA-A8-A08L".

generated words $\{y_i^l\}_{i < n}$ serve as queries, the fused representation \mathbf{F} provides keys and values for the decoder and the n -th word y_n^l is predicted as:

$$y_n^l = \text{CrossAttn}(\{y_i^l\}_{i < n}, \mathbf{F}; \mathbf{\Omega}_l), l \in \{1, 2, \dots, L\}, \quad (3)$$

where $\mathbf{\Omega}_l$ indicates learnable weights for l -th cross-attention layer in the decoder.

3 Experiments and Results

3.1 Implementation Details

I. Datasets: Following [3], we conduct experiments on PathText-BRCA with the same data splits for training, validation, and testing sets. Additionally, we remove duplicate patient samples across these sets, ensuring no patient overlap. Then the training, validation, and test sets contain 796, 88, and 93 samples, respectively.

II. Model Setting: Following [3], WSIs are processed via CLAM [17] to extract non-overlapping 256×256 tissue patches at 10x magnification, the number of encoder and decoder layers are both 3 with 4 attention heads and the embedding size is 512. UNI [4] (pretrained on over 100 million tissue patches and over 100,000

Table 2. Ablation study on key components. WS denotes weight sharing between the visual and knowledge branch. WSL denotes weight sharing across layers in each branch. VTCA/TTCA refer to the visual/textual token cross-attention. W/o VTCA or TTCA refers to conducting self-attention. KR refers to the knowledge retrieval module. "AVG. Δ " represents the average metric promotion compared to the vanilla Transformer.

| WS | WSL | VTCA | KR | TTCA | BLEU-1 | BLEU-2 | BLEU-3 | BLEU-4 | METEOR | ROUGE | Fact _{ent} | AVG. Δ |
|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|---------------------|---------------|
| \times | \times | \times | \times | \times | 0.389 | 0.246 | 0.157 | 0.103 | 0.158 | 0.257 | 0.487 | - |
| \times | \times | \checkmark | \times | \times | 0.436 | 0.283 | 0.186 | 0.126 | 0.174 | 0.287 | 0.499 | 10.84% |
| \times | \checkmark | \checkmark | \times | \times | 0.437 | 0.286 | 0.190 | 0.127 | 0.174 | 0.280 | 0.507 | 11.41% |
| \times | \checkmark | \checkmark | \checkmark | \times | 0.416 | 0.269 | 0.176 | 0.119 | 0.169 | 0.283 | 0.505 | 7.75% |
| \times | \checkmark | \checkmark | \checkmark | \checkmark | 0.441 | 0.288 | 0.190 | 0.130 | 0.178 | 0.290 | 0.515 | 13.06% |
| \checkmark | \checkmark | \checkmark | \checkmark | \checkmark | 0.450 | 0.296 | 0.196 | 0.135 | 0.180 | 0.293 | 0.521 | 15.26% |

WSIs) serves as the visual extractor. Besides, in the knowledge branch, we set $k = 0.4$, $m = 20$, $v = 3$. We employ Adam optimizer with initial learning rate $1e-4$ and weight decay $5e-5$. We adopt beam search with the size of 3 as the sampling method. All experiments is conducted on a single A40-40G GPU.

3.2 Comparison between Ours and Other Methods

We compare with seven image captioning methods in Table 1: two LSTM-based methods: CNN-RNN [25] and att-LSTM [27], five Transformer-based methods: Transformer [24], R2Gen [6] and R2GenCMN [5] (both designed for radiology image captioning), MI-Gen [3] and HistGen [10] (both designed for pathology report generation). Following [3], we adopt four Natural Language Processing (NLP) metrics to evaluate generation performance: BLEU [19], METEOR [1], ROUGE [20] and Fact_{ent} [18]. Besides, we adopt three classification metrics to evaluate Her-2 prediction in generated reports: Precision, Recall and F1 score.

Table 1 shows LSTM-based methods underperform Transformer variants in NLP metrics. R2Gen [6] and R2GenCMN [5] fall short in recognizing key visual features when applied to WSI analysis. MI-Gen [3] feed visual tokens into Transformer to model relationships between patches, which dilutes critical pathological features due to the overwhelming number of redundant regions. Although the local-global hierarchical encoding strategy of HistGen [10] improves efficiency, it still struggles to select key pathological features and lack semantic content. Our historical report guided bi-modal concurrent learning method provides rich semantic content and distills critical pathological features and knowledge features iteratively, achieving 0.135 in BLEU-4 (+14.4% versus HistGen [10]).

Fig. 2 (b) visualizes the attention scores from the 1st layer in the visual token cross-attention module, highlighting tumor tissues versus fibroblast tissues, indicating key feature extraction capability of our method. Fig. 2 (c-e) demonstrate our method effectively generates medical terms consistent with the ground truth, *e.g.*, "m-8500/3", a morphological term in ICD-O-3 [8], indicating invasive ductal carcinoma. Additionally, retrieved knowledge enhances semantics, *e.g.*, "Primary invasive carcinoma of the left breast", improving diagnostic accuracy.




Fig. 3. Performance changes by varying the (a) selecting ratio k , (b) number of knowledge features v and region size m .

3.3 Ablation Study

We conduct ablation studies on five components: 1) Weight Sharing (WS) between the visual and knowledge branch, 2) Weight Sharing across Layers (WSL) in each branch, 3) Visual Token Cross-attention (VTCA), 4) Knowledge Retrieval (KR) and 5) Textual Token Cross-attention (TTCA). As shown in Table 2, the baseline model without any proposed modules (*i.e.*, vanilla Transformer) achieves the lowest performance (Row 1 in Table 2), confirming the necessity of addressing redundancy of WSIs and lack of semantic content. Introducing VTCA (Row 2 in Table 2) and TTCA (Row 5 in Table 2) yield over 10% and 5% relative improvements, representatively. This demonstrates our bi-modal concurrent learning strategy effectively mitigates information redundancy in both WSIs and retrieved knowledge. Enabling WSL (Row 3 in Table 2) further enhances BLEU-4, which suggests weight sharing in each branch stabilizes feature propagation. Simply incorporating the KR module (Row 4 in Table 2) causes performance drop with redundant knowledge. However, by employing TTCA (Row 5 in Table 2), consistent performance gain across all metrics, attributed to recognizing key knowledge embeddings with valuable semantic content. Notably, sharing weights between two branches (Row 6 in Table 2) achieves optimal performance via cross-modal visual-knowledge alignment.

Fig. 3 shows three hyper-parameter sensitivity, with default values $k = 0.4$, $v = 3$ and $m = 20$. As illustrated in Fig. 3 (a), it can be seen that k is not sensitive in $[0.4, 0.8]$. When k is too large, *i.e.*, $k = 1$, which means using all patch embeddings to retrieve knowledge, irrelative patches brings considerable redundant information, resulting in obvious performance degradation. As illustrated in Fig. 3 (b), the best performance can be achieved when v changes from 3 to 7. As illustrated in Fig. 3 (c), performance keeps stable when m varies with the range of $[20, 80]$, which maintains the best performance with lower computation.

4 Conclusion

In this paper, we present a novel historical report guided bi-modal concurrent learning framework for pathology report generation to address two challenges: lack of semantic content and considerable redundancy in WSIs. The proposed knowledge retrieval provides semantically rich diagnostic report information and the bi-modal concurrent learning strategy extracts key pathological features and

key knowledge features. Extensive experiments on the PathText (BRCA) dataset demonstrate state-of-the-art performance both in report generation and Her-2 prediction. This work establishes a new paradigm for pathology report generation which effectively integrates visual evidence and historical knowledge and mitigates information redundancy in both WSIs and retrieved knowledge, with potential applications extending to other medical image interpretation tasks.

Acknowledgements This work was supported by the National Natural Science Foundation of China (Grant No. 62471182), Shanghai Rising-Star Program (Grant No. 24QA2702100), and the Science and Technology Commission of Shanghai Municipality (Grant No. 22DZ2229004).

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

References

1. Banerjee, S., Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization (2005)
2. Chen, L., Zhao, H., Liu, T., Bai, S., Lin, J., Zhou, C., Chang, B.: An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In: ECCV (2024)
3. Chen, P., Li, H., Zhu, C., Zheng, S., Shui, Z., Yang, L.: Wscaption: Multiple instance generation of pathology reports for gigapixel whole-slide images. In: MICCAI (2024)
4. Chen, R.J., Ding, T., Lu, M.Y., Williamson, D.F., Jaume, G., Song, A.H., Chen, B., Zhang, A., Shao, D., Shaban, M., et al.: Towards a general-purpose foundation model for computational pathology. *Nature Medicine* **30** (2024)
5. Chen, Z., Shen, Y., Song, Y., Wan, X.: Cross-modal memory networks for radiology report generation. arXiv preprint arXiv:2204.13258 (2022)
6. Chen, Z., Song, Y., Chang, T.H., Wan, X.: Generating radiology reports via memory-driven transformer. arXiv preprint arXiv:2010.16056 (2020)
7. El Nahhas, O.S., van Treeck, M., Wölflein, G., Unger, M., Ligerio, M., Lenz, T., Wagner, S.J., Hewitt, K.J., Khader, F., Foersch, S., et al.: From whole-slide image to biomarker prediction: end-to-end weakly supervised deep learning in computational pathology. *Nature Protocols* **20** (2025)
8. Fritz, A.G.: International classification of diseases for oncology: ICD-O. World health organization (2000)
9. Gu, T., Liu, D., Li, Z., Cai, W.: Complex organ mask guided radiology report generation. In: WACV (2024)
10. Guo, Z., Ma, J., Xu, Y., Wang, Y., Wang, L., Chen, H.: Histgen: Histopathology report generation via local-global feature encoding and cross-modal context interaction. In: MICCAI (2024)
11. Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T.J., Zou, J.: A visual-language foundation model for pathology image analysis using medical twitter. *Nature medicine* **29** (2023)

12. Li, C., Wong, C., Zhang, S., Usuyama, N., Liu, H., Yang, J., Naumann, T., Poon, H., Gao, J.: Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *NIPS* **36** (2024)
13. Li, H., Zhu, C., Zhang, Y., Sun, Y., Shui, Z., Kuang, W., Zheng, S., Yang, L.: Task-specific fine-tuning via variational information bottleneck for weakly-supervised pathology whole slide image classification. In: *CVPR* (2023)
14. Liu, Q., Louis, P.C., Lu, Y., Jha, A., Zhao, M., Deng, R., Yao, T., Roland, J.T., Yang, H., Zhao, S., et al.: Simtriplet: Simple triplet representation learning with a single gpu. In: *MICCAI* (2021)
15. Lu, J., Xiong, C., Parikh, D., Socher, R.: Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In: *CVPR* (2017)
16. Lu, M.Y., Chen, B., Williamson, D.F., Chen, R.J., Zhao, M., Chow, A.K., Ikemura, K., Kim, A., Pouli, D., Patel, A., et al.: A multimodal generative ai copilot for human pathology. *Nature* **634** (2024)
17. Lu, M.Y., Williamson, D.F., Chen, T.Y., Chen, R.J., Barbieri, M., Mahmood, F.: Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature biomedical engineering* **5** (2021)
18. Miura, Y., Zhang, Y., Tsai, E.B., Langlotz, C.P., Jurafsky, D.: Improving factual completeness and consistency of image-to-text radiology report generation. *arXiv preprint arXiv:2010.10042* (2020)
19. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: *ACL* (2002)
20. Rouge, L.C.: A package for automatic evaluation of summaries. In: *Proceedings of Workshop on Text Summarization of ACL, Spain* (2004)
21. Tan, J.W., Kim, S., Kim, E., Lee, S.H., Ahn, S., Jeong, W.K.: Clinical-grade multi-organ pathology report generation for multi-scale whole slide images via a semantically guided medical text foundation model. In: *MICCAI* (2024)
22. Tran, M., Schmidle, P., Wagner, S.J., Koch, V., Lupperger, V., Feuchtinger, A., Boehner, A., Kaczmarczyk, R., Biedermann, T., Eyerich, K., et al.: Generating highly accurate pathology reports from gigapixel whole slide images with histogpt. *medRxiv* (2024)
23. Tu, T., Azizi, S., Driess, D., Schaeckermann, M., Amin, M., Chang, P.C., Carroll, A., Lau, C., Tanno, R., Ktena, I., et al.: Towards generalist biomedical ai. *NEJM AI* **1** (2024)
24. Vaswani, A.: Attention is all you need. *NIPS* (2017)
25. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: *CVPR* (2015)
26. Xu, H., Usuyama, N., Bagga, J., Zhang, S., Rao, R., Naumann, T., Wong, C., Gero, Z., González, J., Gu, Y., et al.: A whole-slide foundation model for digital pathology from real-world data. *Nature* (2024)
27. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: *ICML* (2015)
28. Zhang, H., Meng, Y., Zhao, Y., Qiao, Y., Yang, X., Coupland, S.E., Zheng, Y.: Dtf-d-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In: *CVPR* (2022)
29. Zhu, L., Wei, F., Lu, Y.: Beyond text: Frozen large language models in visual signal comprehension. In: *CVPR* (2024)