

Clinical-grade Multi-Organ Pathology Report Generation for Multi-scale Whole Slide Images via a Semantically Guided Medical Text Foundation Model

Jing Wei Tan^{*1}, SeungKyu Kim^{*1}, Eunsu Kim², Sung Hak Lee², Sangjeong Ahn³, Won-Ki Jeong¹✉

¹Department of Computer Science and Engineering, College of Informatics,
Korea University, Republic of Korea
{jingwei.92,ksk8804,wkjeong}@korea.ac.kr

²Department of Hospital Pathology, Seoul St. Mary's Hospital, College of Medicine,
The Catholic University of Korea, Republic of Korea
kes311az@gmail.com, hakjjang@catholic.ac.kr

³Department of Pathology, Korea University Anam Hospital, College of Medicine,
Korea University, Republic of Korea
vanitas80@korea.ac.kr

Abstract. Vision language models (VLM) have achieved success in both natural language comprehension and image recognition tasks. However, their use in pathology report generation for whole slide images (WSIs) is still limited due to the huge size of multi-scale WSIs and the high cost of WSI annotation. Moreover, in most of the existing research on pathology report generation, sufficient validation regarding clinical efficacy has not been conducted. Herein, we propose a novel Patient-level Multi-organ Pathology Report Generation (PMPRG) model, which utilizes the multi-scale WSI features from our proposed multi-scale regional vision transformer (MR-ViT) model and their real pathology reports to guide VLM training for accurate pathology report generation. The model then automatically generates a report based on the provided key features-attended regional features. We assessed our model using a WSI dataset consisting of multiple organs, including the colon and kidney. Our model achieved a METEOR score of 0.68, demonstrating the effectiveness of our approach. This model allows pathologists to efficiently generate pathology reports for patients, regardless of the number of WSIs involved. The code is available at <https://github.com/hvcl/Clinical-grade-Pathology-Report-Generation/tree/main>

Keywords: Pathology report generation · Vision language model · Multi-scale whole slide image. · Foundation model

* Equal contribution

✉ Corresponding author: wkjeong@korea.ac.kr

1 Introduction

The generation of pathology reports is vital, as it offers diagnostic insights derived from the analysis of pathology images (a.k.a WSI), establishing a cornerstone for prognosis evaluation, and ultimately improving diagnostic efficiency for clinicians. However, traditional pathology examination requires significant resources, entails labor-intensive procedures, consumes time, and necessitates specialized expertise. It may take up to several days to scan and examine slides and generate a pathology report. Therefore, there exists an increasing demand for automating the report generation process.

VLMs have gained traction in pathology report generation (PRG). Prior works like Zhang *et al.* [27] combined CNN and LSTM for VLM training using ROI or sampled patches from WSIs. PLIP [7], MI-Zero [12] and CITE [26] adopted CLIP-based training of patch-text pairs, facilitating zero-shot transfer. GPC [14] merged CNN with a Transformer-based language model for learning from diverse pathology patches and text pairs. However, these studies predominantly focus on patch and textual data, requiring specialized expertise for defining ROIs, which is laborious. Recent approaches in WSI-level report generation include models like Zhang *et al.* [25], which used CNN and LSTM with thumbnail images and a bag of instances to represent entire WSIs. Sengupta *et al.* [19] and Guevara *et al.* [6] utilized pretrained HIPT [5] for WSI feature extraction and trained with CLIP or LSTM. MI-gen [4] introduced to insert a Position-Aware Module which adds positional information to each instance into the output obtained through the transformer encoder block. However, current approaches are limited to generating simple, unstructured descriptions of WSI, often used for biopsy samples, and its clinical efficacy is not yet validated.

The main motivation of our work stems from the following observations. Firstly, unlike the reports generated by most existing work, real clinical reports exhibit less descriptive but well-structured formats, varying based on the target organ and disease (see an example report in Fig. 1(A).) Existing methods [25, 19, 6] often rely on highly compressed 1D representations from WSIs, raising concerns about their clinical efficacy given the diversity of pathology reports. As we extend to multi-organ settings, the variability in disease topics intensifies, further emphasizing the importance of addressing this issue. Secondly, pathologists commonly analyze multiple slides per patient to ensure diagnostic accuracy, necessitating an automatic PRG based on multiple WSIs. Therefore, developing an automatic PRG based on multiple WSIs from a single patient is more closely aligned with clinical practice. In such a case, the weak-label problem, a commonly found issue with WSI, arises from the predominant assignment of labels at the slide or patient level, and multiple instance learning (MIL) is commonly employed to address this problem (such as DSMIL [10] and ZoomMIL [23]). However, we observed that conventional MIL approaches often prioritize class-specific tasks, such as cancer sub-type classification. This emphasis poses a challenge when directly applying MIL to PRG tasks, which encompass a wide array of diagnostic information. Moreover, we also observed a need to develop an efficient and lightweight encoder that leverages multi-scale information from WSI.

We empirically found that existing multi-scale WSI encoders are not efficient for training using a large number of WSIs due to the architecture design.

In this work, in contrast to traditional methods relying on patches and text, we leverage the multi-scale nature of WSIs along with actual pathology reports to train a VLM tailored for PRG. Inspired by HIPT [5], which learns features from multiple FoVs on WSIs, we introduce MR-ViT to capture representative features of each multi-scale region. With these regional features, we also propose a Patient-level Multi-organ Pathology Report Generation (PMPRG) model using a medical text foundation model to handle the patients with varying numbers of WSIs, each of different sizes and various scales, making it more suitable for real-world scenarios. The main contributions of our work can be summarized as follows:

- We propose an efficient unsupervised regional hierarchical model for WSIs, offering improved efficiency in multi-scale feature representation. By integrating multiple instances from diverse FoVs or multi-scales through a simple two-level feature encoding process, our model captures both local details and broader context effectively. This enables the learning of patient-level features alongside text within our model and allows for easy fine-tuning of our report-generation model without compromising quality. Experimental results validate the superiority of our approach over SOTA multi-scale methods.
- We propose an automated pathology report generation model that efficiently handles WSIs with various report information while allowing easy extension to multiple organs. This approach is suitable for real-world scenarios where organ attributes vary significantly, and it offers the advantage of generating clinically valid reports at both patient-level and WSI-level.
- To the best of our knowledge, this work is the first attempt to create an explainable model that can generate clinical-grade pathology reports, trained using in-house multi-organ datasets consisting of 7422 WSIs with clinical reports in inhomogeneous formats and structures.

2 Method

2.1 Overview and Pre-processing

Our proposed model mainly composed of two stages: (1) train a multi-scale regional feature extractor in an unsupervised manner and (2) train a report generator with the regional feature from the previous stage with reports and labels extracted from reports as shown in Fig. 1(C) and (A), respectively.

The input WSIs are first pre-processed by extracting the patches from $1.25\times$ and $5\times$ magnification level in $256\times 256\times 3$ dimension. Our dataset includes patients with multiple WSIs, each accompanied by a single pathology report. These reports contain various descriptions for each disease topic (Fig.1(A)), varying significantly across organs but showing similarity within the same organ. Inspired by radiology report generation methods [8, 22, 24, 11, 1], we established predetermined key-tags for each organ in reports, relying exclusively on WSIs guided by




Fig. 1. Overview of our proposed pipeline: (A) PMPRG, (B) Multi-scale WSI, (C) MR-ViT, (D) Tag-guided feature extractor and (E) GPT-2.

pathologists for accurate predictions. Sentences containing these tags were extracted using a rule-based method. If key-tag-related items are missing, they are excluded from loss calculation. We assign tag labels based on diagnosis scenarios after preprocessing, and also each report is labeled with an organ class. The output includes three labels for each report: sentences $Y_{sen} = \{y_{sen_i}\}_{i=1}^K$, tag class labels $Y_{tag} = \{y_{tag_i}\}_{i=1}^K$, and organ class labels $Y_{org} = \{o \mid o \in \{0, 1, \dots, n_o\}\}$, where K and n_o denote number of tags specific to each organ and total number of organ types within the dataset, respectively. The value of K may vary based on the type of organ.

2.2 Multi-scale Regional Vision Transformer (MR-ViT)

When adapting HIPT [5] for multi-scale regional representation learning, we introduce two key distinctions from the original approach: (1) utilizing patches of different scales from a specific region as the model input, and (2) reducing computational resources by requiring only two stages to obtain the WSI representation. In the initial phase of MR-ViT_R (Fig. 1(C)(i)), we construct a 1024×1024 dimensional image representing a specific region r by selecting one patch at level 2 and randomly choosing 15 out of 16 patches at level 1 as shown in Fig. 1(B). Each region consists of 16 visual tokens (256×256), serving as input for ViT. Subsequently, we extract features for these visual tokens using a pre-trained model C such as VGG16. These extracted embeddings F_C are then utilized in DINO [3] to capture multi-scale features for each specific region. Following this, the MR-ViT_R model is employed to extract multi-scale regional representations F_R for each WSI. F_R are then utilized for training MR-ViT_S, as depicted in Fig. 1(C)(ii) later on. We randomly select Q (at $L \times L$ dimension)

regions from F_R and get F'_R , which are then fed into the second stage MR-ViT_S. At this stage, we aim to capture the multi-scale regional features representing each WSI well in a single vector. Subsequently, we obtain a WSI representation $F_S = \text{MR-ViT}_S(F'_R)$. The DINO model hyper-parameter setting is exactly same as HIPT, except for the input size and output size.

2.3 Pathology Report Generation

As illustrated in Fig. 1(A), our model at this stage comprises 6 primary modules: **1)** MR-ViT, **2)** Organ Classifier (Cls_{org}), **3)** Tag-Selection Dictionary, **4)** Tag-specific Feature Extractor ($F_{TE}(\cdot)$), **5)** Tag Classifier (Cls_{tag}), and **6)** Conditional Language Model ($T(\cdot)$).

MR-ViT. All WSIs corresponding to a single patient are patched and processed through the MR-ViT module, generating multiple regional representations. These representations from multiple WSIs for each patient are then concatenated to form $F_{R,pat}$.

Organ Prediction. During batch training, padding is applied along the sequence dimension to $F_{R,pat}$, resulting $F_{R,pat}^*$. Subsequently, after applying average pooling, $F_{R,pat}^*$ is being supplied to the organ-classifier (Cls_{org}), consisting of 2 sequential linear layers. The organ prediction loss is formulated as follows:

$$L_{org} = CELoss(logits_{org}, Y_{org})$$

Tag-Selection Dictionary. Given that we have predefined key-tags for each organ, the tags remain consistent for the same organ. Initially, we internally define learnable tag embeddings within the model, denoted as $tags_j = \{tag_i\}_{i=1}^K$, and $tag_{tot} = \{tags_j\}_{j=1}^{n_o}$. Subsequently, based on the results of organ prediction from the previous module, the corresponding organ's tag-embedding set is retrieved as below:

$$tag_batch = \{tags_{j=Pred_{org}[k]}\}_{k=1}^B$$

where $Pred_{org} = Argmax(logits_{org})$.

Tag-specific Feature Extractor. Tag-specific Feature Extractor ($F_{TE}(\cdot)$) comprised of *Decoder Block* ($G_d(\cdot)$) and *Attention Pooler* ($H(\cdot)$), which takes tag_batch and $F_{R,pat}^{**}$ as inputs for extracting tag-attended visual features (f_{vt}).

$$\begin{aligned} f_{vt} &= F_{TE}(tag_batch, F_{R,pat}^{**}) \in R^{B \times K \times d}, \\ &= H(G_d(tag_batch, F_{R,pat}^{**}), tag_batch) \end{aligned}$$

where $F_{R,pat}^{**}$ represents linear projection of $F_{R,pat}^*$, aligning dimensions with the tags. $F_{TE}(\cdot)$ block follows a standard transformer decoder structure, comprising self-attention, cross-attention, and feed-forward layers. Here, $F_{R,pat}^{**}$ is set as the query and tag_batch as the key and value. Similarly, $H(\cdot)$ is composed of cross-attention and feed-forward layers, but does not include a self-attention layer. In

this case, the query is set as tag_batch and the key and value are the final output of $G_d(\cdot)$. After obtaining f_{vt} , we reshape it to $f'_{vt} \in R^{(B \times K) \times d}$ and pass it to Cls_{tag} which has similar structure to Cls_{org} , enabling it to make predictions for each tag. The formulation for tag classifier loss is as follows:

$$L_{tag} = CELoss(logits_{tag}, Y_{tag})$$

where $logits_{tag}$ represents the tag label prediction.

Conditional Language Generation. By feeding the obtained visual features (f'_{vt}) into the language model ($T(\cdot)$), we enable the model to generate descriptions corresponding to each tag. For this task, we employed a pre-trained GPT-2 model [17] sourced from PubMed [15]. By substituting the self-attention layers of GPT-2 with pseudo self-attention (PSA) [13] layers, we conditioned the language model accordingly. Our approach closely follows the methodologies outlined in previous studies [2, 26]. To preserve the strengths of the foundational language model and for resource-efficient training, we exclusively trained the projection parameters of each block for conditioning input while keeping the remaining parameters frozen. Consequently, we derive logits $logits_{sen} = T(f'_{vt}, f_i)$ for each sentence in the report.

$$L_{sen} = CELoss(logits_{sen}, Y_{sen})$$

Here, f_i represents the token embedding of the report. During training, right-shifted Y_{sen} is used as input, and during inference, only the $< BOS >$ token is used. Finally, the model is trained based on 3 different loss components.

$$L_{tot} = \alpha L_{org} + \beta L_{tag} + \gamma L_{sen}$$

where $\alpha=0.2$, $\beta=0.6$, and $\gamma=0.2$ are weights assigned to each loss component.

3 Result

In this study, we collected data from Korea University Anam Hospital covering two organs: the kidney and colon. The dataset includes 1991 patients, with each patient having multiple WSIs, totaling 7422 slides overall. We partitioned the dataset randomly into three subsets: 70% for training, 20% for validation, and

Table 1. The performance comparison of patient-level multi-scale image encoder with the SOTA methods of multi-scale WSI and an ablation study of different pre-trained models used in MR-ViT.

Method	Accuracy	
	Diagnosis Type	Tumor Grade
HIPT [5]	0.5455	0.3030
ZoomMIL [23]	0.8169	0.4688
MR-ViT (Tan <i>et al.</i> [21])	0.7273	0.3939
MR-ViT (BT [9])	0.7879	0.3939
MR-ViT (VGG16 [20])	0.8485	0.4242

10% for testing. For PRG evaluation metrics, in addition to the commonly used Natural Language Generation (NLG) metrics, we integrated Clinical Efficacy (CE) metrics to evaluate the model’s clinical effectiveness in making predictions. For NLG assessment, we utilized three well-established measures: BLEU-n [16], METEOR [2], and ROUGE-L [18]. We measure F1 score and total accuracy for the CE metric. Implementation details and additional dataset description are provided in the supplementary materials.

3.1 Experimental Result

Image Encoder. In this study, we assessed the performance of our proposed MR-ViT by extracting WSI features and classifying them into two specific tasks: (1) diagnosis type and (2) tumor grade. Given our focus on patient-level rather than slide-level analysis, we obtained slide-level predictions and used the maximum value among the predicted slide classes as the final prediction. We compared our approach with two state-of-the-art (SOTA) multi-resolution WSI methods, HIPT [5] and ZoomMIL [23]. Despite attempts to fine-tune the HIPT model for regional feature extraction with our dataset, we encountered significant challenges, including HIPT requiring approximately 300 times longer training time for a single WSI compared to our model under equivalent computational resources. Consequently, we utilized the pre-trained HIPT model directly for feature extraction as in previous studies [19, 6], while training ZoomMIL from scratch in a task-specific manner. In our experiments (Table 1), MR-ViT with the pre-trained VGG model achieved superior performance compared to both HIPT and ZoomMIL for diagnosis type classification. However, ZoomMIL outperformed our MR-ViT and HIPT models for tumor grade classification, likely due to its supervised end-to-end MIL approach whereas MR-ViT and HIPT are trained in unsupervised manner. Additionally, we utilized pre-trained encoders such as a pathology image BT model [9], a multi-scale contrastive encoder [21], and the ImageNet VGG16 model [20] to extract patch features for MR-ViT_R.

Table 2. Performance of report generation. For accuracy, we average the individual score of each class (total 26 classes). B-n, M, and R-L refers to BLEU, METEOR, and Rouge-L, respectively. Methods with * sign indicates generating whole report at once.

Method	CE metrics		NLG metrics					
	F1	Acc.	B-1	B-2	B-3	B-4	M	R-L
HIPT _S *	-	-	0.1456	0.1387	0.1305	0.1187	0.2311	0.2880
MR-ViT _S *	-	-	0.3952	0.3739	0.3518	0.3319	0.4797	0.5137
HIPT _S	0.1808	0.2212	0.4053	0.3589	0.3298	0.3046	0.5381	0.4501
MR-ViT _S	0.4384	0.4285	0.4948	0.4566	0.4274	0.3979	0.6334	0.5449
HIPT _R	0.1392	0.1913	0.3866	0.3364	0.3057	0.2819	0.5122	0.4194
w/o G_d , Cl_{stag}	-	-	0.5292	0.4982	0.4739	0.4490	0.6670	0.5888
w/o G_d	0.5786	0.5886	0.5362	0.5029	0.4771	0.4500	0.6813	0.5922
w/o Cl_{stag}	-	-	0.5066	0.4719	0.4444	0.4154	0.6570	0.5747
MR-ViT _R (Ours)	<u>0.5773</u>	0.6022	0.5507	0.5184	0.4925	0.4654	0.6834	0.6033




Fig. 2. (a) The attention map depicts the importance of WSI regions for specific tags, where brighter regions indicate higher importance. (b) Examples of reports generated using our method and baseline methods.

The results indicated that MR-ViT with the pre-trained VGG model showed better performance than the other pre-trained models.

PMPRG. We devised three main comparison scenarios for our model: 1) representative visual features extracted using only visual encoder are directly fed into the language model, generating the entire report in one step. For patients with multiple WSIs, we utilized the average-pooled features. 2) follow our overall pipeline but utilizes slide-level features extracted without semantic guidance, and 3) follow the approach we proposed. In each scenario, we compared our encoder (MR-ViT) with HIPT. Due to computational resource constraints, HIPT [5] was utilized with pre-trained weights without further training, following default settings. Table 2 presents comprehensive results for report generation, highlighting the effectiveness of segmenting the report into distinct tags and transforming visual features into specific vectors for each tag, compared to directly relaying features from the visual encoder to the language model. This approach, especially for HIPT, results in a notable improvement in METEOR scores (approximately 0.3 increase from HIPT_S* to HIPT_S), enhancing report generation quality. Moreover, providing semantic guidance for each tag during the feature compression stage significantly contributes to generating high-quality reports compared to using slide-level features extracted without guidance (MR-ViT_S → MR-ViT_R), notably improving CE metric performance (approximately 17% more accurate predictions). Although HIPT exhibits a decrease in both CE and NLG metrics, it is speculated that ineffective feature extraction from the beginning may have contributed to limited impact on the results, as evidenced by notably low scores achieved by HIPT. Additionally, an ablation study for each module reveals that

each module contributes to guiding the model’s report generation favorably. Specifically, the Cls_{tag} module aids in generating clinically accurate descriptions for each tag, resulting in improvements of 0.03 in both METEOR and Rouge-L scores, along with an increase of approximately 0.05 in BLEU scores.

Fig. 2(a) illustrates qualitative results with attention maps for interpreting the model’s predictions. WSIs displayed are from a single patient, with the model integrating information from multiple WSIs for a comprehensive diagnosis. Attention weights are obtained from $H(\cdot)$ and yield individual attention maps for each tag, revealing where the model focuses its attention and providing insights into specific regions considered during diagnosis. Fig. 2(b) showcases the generated reports from each model. Generating the entire report at once (HIPT *_S , MR-ViT *_S) may lack descriptions for tags present in the actual report. Conversely, our model includes all tags, which seems reasonable considering the consistent nature of pathology reports within the same organ. More results are provided in supplementary material.

4 Conclusion

In this study, we present a streamlined multi-scale regional encoder for WSIs and an innovative patient-level multi-organ pathology generation model. Our proposed image encoder enhances training efficiency for multi-scale WSI images compared to the baseline method. Moreover, by harnessing the multi-scale regional features, our pathology report generation model can provide pathologists with practical, real-world pathology reports. Future plans include extending our model to incorporate three magnification scales and incorporating more data from diverse organs.

Acknowledgements. This study was approved by the institutional review board of Korea University Anam Hospital (IRB NO.2024AN0190). This work was partially supported by the National Research Foundation of Korea (RS-2024-00349697, NRF-2021R1A6A1A13044830), the Institute for Information & Communications Technology Planning & Evaluation (IITP-2024-2020-0-01819), the Technology development Program(RS-2024-00437796) funded by the Ministry of SMEs and Startups(MSS, Korea), the Korea Health Industry Development Institute (RS-2021-KH113146) and a Korea University Grant.

Disclosure of Interests. There are no conflicts of interest to declare.

References

1. Alfarghaly, O., Khaled, R., Elkorany, A., Helal, M., Fahmy, A.: Automated radiology report generation using conditioned transformers. *Informatics in Medicine Unlocked* **24**, 100557 (2021)

2. 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. pp. 65–72 (2005)
3. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
4. Chen, P., Li, H., Zhu, C., Zheng, S., Yang, L.: Mi-gen: Multiple instance generation of pathology reports for gigapixel whole-slide images (2023)
5. Chen, R.J., Chen, C., Li, Y., Chen, T.Y., Trister, A.D., Krishnan, R.G., Mahmood, F.: Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16144–16155 (2022)
6. Guevara, B.C., Marini, N., Marchesin, S., Aswolinskiy, W., Schlimbach, R.J., Podareanu, D., Ciompi, F.: Caption generation from histopathology whole-slide images using pre-trained transformers. In: Medical Imaging with Deep Learning, short paper track (2023)
7. 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**(9), 2307–2316 (2023)
8. Jing, B., Xie, P., Xing, E.: On the automatic generation of medical imaging reports. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/p18-1240>, <http://dx.doi.org/10.18653/v1/P18-1240>
9. Kang, M., Song, H., Park, S., Yoo, D., Pereira, S.: Benchmarking self-supervised learning on diverse pathology datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3344–3354 (2023)
10. Li, B., Li, Y., Eliceiri, K.W.: Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14318–14328 (2021)
11. Liu, F., Wu, X., Ge, S., Fan, W., Zou, Y.: Exploring and distilling posterior and prior knowledge for radiology report generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13753–13762 (2021)
12. Lu, M.Y., Chen, B., Zhang, A., Williamson, D.F., Chen, R.J., Ding, T., Le, L.P., Chuang, Y.S., Mahmood, F.: Visual language pretrained multiple instance zero-shot transfer for histopathology images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19764–19775 (2023)
13. Luke, M.K., Sebastian, G., Rush Alexander, M., et al.: Encoder-agnostic adaptation for conditional language generation. *arXiv* (2019)
14. Nguyen, A.T., Kwak, J.T.: Gpc: Generative and general pathology image classifier. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 203–212. Springer (2023)
15. Papanikolaou, Y., Pierleoni, A.: Dare: Data augmented relation extraction with gpt-2 (2020)
16. Papineni, K., Roukos, S., Ward, T., Zhu, W.: A method for automatic evaluation of machine translation”. the Proceedings of ACL-2002, ACL, Philadelphia, PA, July 2002 (2001)
17. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners (2019)

18. ROUGE, L.C.: A package for automatic evaluation of summaries. In: Proceedings of Workshop on Text Summarization of ACL, Spain. vol. 5 (2004)
19. Sengupta, S., Brown, D.E.: Automatic report generation for histopathology images using pre-trained vision transformers. arXiv preprint arXiv:2311.06176 (2023)
20. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
21. Tan, J.W., Nguyen, K.T., Lee, K., Jeong, W.K.: Multi-scale contrastive learning with attention for histopathology image classification. In: Medical Imaging 2023: Digital and Computational Pathology. vol. 12471, pp. 294–301. SPIE (2023)
22. Tanida, T., Müller, P., Kaissis, G., Rueckert, D.: Interactive and explainable region-guided radiology report generation. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (Jun 2023). <https://doi.org/10.1109/cvpr52729.2023.00718>, <http://dx.doi.org/10.1109/CVPR52729.2023.00718>
23. Thandiackal, K., Chen, B., Pati, P., Jaume, G., Williamson, D.F., Gabrani, M., Goksel, O.: Differentiable zooming for multiple instance learning on whole-slide images. In: European Conference on Computer Vision. pp. 699–715. Springer (2022)
24. You, D., Liu, F., Ge, S., Xie, X., Zhang, J., Wu, X.: Aligntransformer: Hierarchical alignment of visual regions and disease tags for medical report generation (2022)
25. Zhang, R., Weber, C., Grossman, R., Khan, A.A.: Evaluating and interpreting caption prediction for histopathology images. In: Machine Learning for Healthcare Conference. pp. 418–435. PMLR (2020)
26. Zhang, Y., Gao, J., Zhou, M., Wang, X., Qiao, Y., Zhang, S., Wang, D.: Text-guided foundation model adaptation for pathological image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 272–282. Springer (2023)
27. Zhang, Z., Chen, P., McGough, M., Xing, F., Wang, C., Bui, M., Xie, Y., Sapkota, M., Cui, L., Dhillon, J., et al.: Pathologist-level interpretable whole-slide cancer diagnosis with deep learning. *Nature Machine Intelligence* **1**(5), 236–245 (2019)