

Leveraging Image Captions for Selective Whole Slide Image Annotation

Jingna Qiu¹, Marc Aubreville^{2,3}, Frauke Wilm^{1,4}, Mathias Öttl^{1,4}, Jonas Utz¹,
Maja Schlereth¹, and Katharina Breininger^{1,5}

¹ Department Artificial Intelligence in Biomedical Engineering,
FAU Erlangen-Nürnberg, Erlangen, Germany jingna.qiu@fau.de

² Technische Hochschule Ingolstadt, Ingolstadt, Germany

³ Flensburg University of Applied Sciences, Flensburg, Germany

⁴ Pattern Recognition Lab, Department of Computer Science,
FAU Erlangen-Nürnberg, Erlangen, Germany

⁵ Center for AI and Data Science, Universität Würzburg, Würzburg, Germany

Abstract. Acquiring annotations for whole slide images (WSIs)-based deep learning tasks, such as creating tissue segmentation masks or detecting mitotic figures, is a laborious process due to the extensive image size and the significant manual work involved in the annotation. This paper focuses on identifying and annotating specific image regions that optimize model training, given a limited annotation budget. While random sampling helps capture data variance by collecting annotation regions throughout the WSI, insufficient data curation may result in an inadequate representation of minority classes. Recent studies proposed diversity sampling to select a set of regions that maximally represent unique characteristics of the WSIs. This is done by pretraining on unlabeled data through self-supervised learning and then clustering all regions in the latent space. However, establishing the optimal number of clusters can be difficult and not all clusters are task-relevant. This paper presents prototype sampling, a new method for annotation region selection. It discovers regions exhibiting typical characteristics of each task-specific class. The process entails recognizing class prototypes from extensive histopathology image-caption databases and detecting unlabeled image regions that resemble these prototypes. Our results show that prototype sampling is more effective than random and diversity sampling in identifying annotation regions with valuable training information, resulting in improved model performance in semantic segmentation and mitotic figure detection tasks. Code is available at <https://github.com/DeepMicroscopy/Prototype-sampling>.

Keywords: Whole slide images · histopathology · annotation region selection · low-data learning · image caption.

1 Introduction

Deep learning models for whole slide image (WSI) semantic segmentation and mitotic figure detection can assist in tumor grading by identifying possible ma-

lignant areas [14,14,26]. High-quality annotations are crucial to properly train these models, but the acquisition process of these annotations requires pathologists as annotators and is extremely time-consuming due to the large image size in the gigapixel range and the substantial manual effort involved in the annotation. Segmentation masks are generated at the pixel level, and careful examination of thousands of cells is necessary to exhaustively identify mitotic figures (cells undergoing division). This work focuses on discovering specific annotation regions that optimize model training, while leaving the rest of the image unlabeled. This can help reduce the need for manual labeling or accommodate a limited availability of expert annotators. Decreasing the annotation area has been shown to efficiently streamline the process of producing segmentation masks for WSIs [13,15,22,27]; but, to our knowledge, it has not been used for acquiring mitotic figure annotations. Here, annotators are frequently asked to exhaustively annotate manually selected regions of interest [3], which could be considerably simplified if cell examination was required only for a fraction of these regions that are identified as informative.

Given a pool of unlabeled WSIs and a certain annotation budget, the task is to find annotation regions that contain useful information for training the model in the downstream task. One natural idea is random sampling, which ensures data diversity by including tissue from different parts of the WSI. However, it may raise issues such as the selected regions inadequately representing minority classes. Some works involve manually selecting annotation regions [13], which necessitates expertise in estimating which information is most useful for model training. However, humans may not fully recognize the varying degrees of visual variety presented within each class [25]. In recent years, self-supervised learning (SSL) techniques have popularized diversity sampling, which identifies a set of regions that represent distinct features of the dataset with minimum duplication. Specifically, the pretrained model creates embeddings for all regions, the embeddings are clustered and annotation regions are selected at the center of each cluster [10,21] or to maximally cover each cluster [12,28]. However, the optimal number of clusters is often not clear, and some clusters may feature irrelevant information for training the downstream task model, such as artifacts.

In this work, we introduce a novel method called *prototype sampling* for the specified task. It identifies annotation regions that exhibit typical characteristics of each task-specific class without the need for clustering. Specifically, our method involves collecting prototype embeddings for each class from extensive histopathology image-caption databases, and identifying unlabeled regions that closely resemble these prototypes. Two databases are used in the study: *ARCH*, composed of figures and captions from PubMed articles and textbooks [9], and *OpenPath*, encompassing Tweets where pathologists discussed cases [11]. These resources significantly decrease the time and challenge of searching for class prototypes. They offer a valuable and diverse collection of images showcasing both typical and ambiguous aspects of many diseases and include a wide range of data variations, including staining material and case origin. Similar to diversity sampling, our approach uses a pretrained model to extract feature embeddings for

the following sampling process. However, we assign unlabeled regions to task-specific classes based on detected prototypes instead of relying on clustering results. One use case of our method is the construction of the initial labeled set of region-based active learning (AL) [19], which iteratively trains the model on existing labeled regions and uses the trained model to determine which regions to annotate next. An initial labeled set that efficiently trains the first AL model can more reliably identify informative annotation regions in the subsequent AL cycle. Another application is to create a preliminary model by annotating a small portion of the data, and then use the model to automate further annotations or propose reference annotations [5]. We test our method on semantic segmentation and mitotic figure detection tasks by training the model using the annotation regions selected through random, diversity, and prototype sampling.

2 Method

2.1 General Setup

Given a pool of N unlabeled WSIs and an annotation budget of n regions of size $l \times l$ per WSI, the objective is to identify annotation regions containing useful training information for the downstream task. The annotation regions are detected based on a specified sampling method. In this section, we outline random sampling and diversity sampling, which serve as benchmark methods for our experiments. The random sampling method selects n regions from each WSI at random locations; regions with less than 10% tissue area are excluded by tissue detection [20]. We mainly followed [10] for diversity sampling. First, each WSI is partitioned into a grid of regions of size $l \times l$. For each region we calculate the feature embedding using a pretrained model. We then aggregate the embeddings of all regions from N WSIs and allocate them to $N * n$ clusters with K-means clustering, in order to prevent the selection of similar images regions from separate WSIs. Finally the regions centered at each cluster are selected. Next we describe prototype sampling in detail.

2.2 Prototype Sampling

Prototype sampling has two main components: 1) the acquisition of prototypes of each task-specific class from image-caption databases by keyword search or text-to-image retrieval, and 2) the construction of a similarity map that shows the similarity of different image areas to the prototypes. The similarity map is used for the identification of annotation regions. An illustration is shown in Fig. 1.

Class Prototypes Identification We obtain typical images of a specific class from histopathology image-caption databases by *keyword search* or *text-to-image retrieval*. *Keyword search* exploits the detailed information contained in captions and is straightforward to conduct (Fig. 1 (1a)). Specifically, we create a set of “with” and “without” keywords, and recognize images with captions that include all “with” keywords and none of the “without” keywords. The “with” keywords




Fig. 1: Workflow of prototype sampling.

are created using class names with several synonyms (e.g., “cancer” and “tumor”), under the assumption that data analysis experts lack medical expertise. The “without” keywords are used to exclude irrelevant images, such as those of type photomicrograph. The recognized images are encoded into embeddings $\mathcal{F}_{prototype}^{(c)}$ using a pre-trained image encoder, with c representing the class. *Text-to-image retrieval* employs text and image encoders created through contrastive language-image pretraining (CLIP) [23] to align visual concepts from images with the perception contained in the text (Fig. 1 (1b)). We use a text prompt “An H&E image of {class_name} tissue.” and obtain its embedding from the text encoder. All images in the database are encoded using the image encoder to create $\mathcal{F}_{database}$. $N_{prototype}$ database image embeddings that have the highest cosine similarity to the text prompt embedding are selected to form $\mathcal{F}_{prototype}^{(c)}$.

Similarity Map Construction and Annotation Region Selection We now identify image regions from the WSI that are representative of the specified class. We partition the WSI $X_i \in \mathbb{R}^{3 \times H_i \times W_i}$ into patches with a stride of s and then calculate their embeddings individually using a pretrained model. We denote a patch as $x \in \mathbb{R}^{3 \times d \times d}$ and its embedding as $f_x \in \mathbb{R}^{C_f}$, where C_f denotes the embedding dimension. The size of the patches d and the level of magnification used to extract the patches from the WSI pyramid are chosen according to the task. For instance, a small patch extracted at high magnification is preferable for comparing its resemblance to prototype images of mitotic figures, whereas patches characterizing a gland should be sufficiently large to encompass all tissue components. Note that each annotation region of size l may contain multiple patches of size d . The representativeness of patch x for class c is defined as the

highest cosine similarity between f_x and the embeddings in $\mathcal{F}_{prototype}^{(c)}$, calculated as $\text{sim}(f_x, \mathcal{F}_{prototype}^{(c)}) = \max_{f_p \in \mathcal{F}_{prototype}^{(c)}} \text{sim}(f_x, f_p)$. Patch similarities are then combined into a similarity map $M_i \in [0, 1]^{\frac{H_i}{s} \times \frac{W_i}{s}}$ for the WSI. The map is used to select annotation regions using a standard or an adaptive region selection method, denoted as “prototype (standard)” and “prototype (adaptive)” accordingly. The standard method [19] moves a sliding window of the annotation region size (i.e., $\frac{l}{s} \times \frac{l}{s}$ (note that l is defined at the original resolution)) across M_i with a stride of 1 pixel, evaluates the similarity of each region by summing up the similarities of patches within it, and then use non-maximum suppression to identify non-overlapping regions with the highest similarity values. The adaptive method [22] allows for selecting regions that can adjust in shape and size to accommodate variations in histopathological structures. The process includes the following steps: 1) Locate the pixel with the highest similarity (u, v) ; 2) Create a binary mask by using a similarity threshold; 3) Identify the connected component that includes (u, v) and choose its bounding box; and 4) Determine the similarity threshold through bisection search to ensure that the bounding box falls within the range $[\frac{1}{2}l \times \frac{1}{2}l, \frac{3}{2}l \times \frac{3}{2}l]$.

3 Experiments

3.1 Image-caption Pair Databases

Two histopathology image-caption databases are used in the study. The ARCH database [9] contains 8,617 histology or immunohistochemistry image-caption pairs extracted from PubMed medical articles and 3,199 pairs from 10 textbooks. The captions are manually curated to only include text related to diagnostic and morphological descriptions. The OpenPath database [11] consists of 116,504 image-caption pairs from Twitter posts and 59,869 pairs from replies, gathered from 32 pathology subspecialty-specific hashtags. The database does not contain the extracted images and captions but only the links to the original posts. Image and caption embeddings obtained from PLIP [11], which is trained by fine-tuning a pre-trained CLIP model [23] on OpenPath, are included in the database.

3.2 Semantic Segmentation

Dataset The public dataset CAMELYON16 [17] consists of 399 WSIs of Hematoxylin & Eosin (H&E)-stained excised regional sentinel lymph nodes. The task is the detection of the presence and extent of breast cancer metastases. Each WSI with metastases is accompanied by a segmentation mask that outlines all metastases at the pixel level. We followed data usages in [22] for fair comparisons. **Training Setups** We followed [22] using a segmentation framework based on patch classification, including all training setups, except for replacing the MobileNet [24] encoder with the two following encoders: 1) *ResNet18_SSL* [7]: A ResNet18 that is self-supervised pretrained on 0.4 million histological patches with SimCLR [6]; 2) *ViT_PLIP*: The image encoder (ViT-32-Base) of PLIP.

Training and inference patches were extracted at the resolution of $0.25 \frac{\mu\text{m}}{\text{px}}$ with 256×256 pixels and 224×224 pixels for the two encoders, respectively.

Creating Class Prototypes We used the ARCH and OpenPath databases for obtaining class prototypes for the experiments with the two encoders. From ARCH, we extracted breast cancer images by keyword search. We used {breast, abnormal (tumor, cancer, carcinoma, metastases, metastasis, metastatic)} and {IHC (immunohistochemical, immunohistochemistry, immunostain), photomicrograph (photomicrography)} as the “with” and “without” keywords, respectively, where synonyms are given in the parentheses. In total 21 images were found. Details on the pre- and post-processing steps can be found in Table S1. These images were processed by the *ResNet_SSL* encoder to generate image embeddings. Prior to this, the images were center-cropped to the size of model input to maintain the original magnification and prevent distortion of the structures. From the OpenPath database, where the image embeddings are provided, we selected the top $N_{\text{prototype}} = 100$ image embeddings under the hashtag “#Breast-Path” that best align with the embedding of the context prompt of “An *H&E* image of breast tumor tissue.” generated by the PILP text encoder.

WSI Similarity Map Calculation Patches of size 256×256 pixels were extracted at resolution $0.25 \frac{\mu\text{m}}{\text{px}}$ with a stride of 256 pixels to ensure the incorporation of details required for metastases identification.

Comparison Methods We compared to random sampling and diversity sampling. For fair comparisons, the two encoders were used to obtain region features in diversity sampling instead of pretraining a new model on the unlabeled data. Prior to this, the regions were downsampled to match the model input sizes.

Evaluation Scenarios and Metrics Recent works have found that the efficiency of a sampling strategy may differ depending on the the number and size of the annotation regions [18,19,22], we follow [22] to evaluate on nine hyperparameter settings by combining $n \in \{1, 3, 5\}$ and $l \in \{4096, 8192, 12288\}$ pixels. We denote these three sizes as S (small), M (medium) and L (large), respectively. Slide-averaged intersection over union across test slides containing tumor (mIoU (tumor)) was used as the evaluation metric, following [22].

3.3 Mitotic Figure Detection

Dataset The MITOS_WSI_CMC [2] dataset consists of 21 slides of canine mammary carcinoma that are fully annotated with 13,907 mitotic figures. The task is to identify all mitotic figures in the WSIs, as mitotic count is a key criteria for tumor grading [8]. We followed data usages in [2] for fair comparison.

Training Setups We followed [2] using a RetinaNet [16] (ResNet18 backbone pretrained on ImageNet) as the mitotic figure detection model and the same training schemes, the only modification was sampling 1,000 training patches per epoch instead of 5,000 when the annotated area was $< 5\%$ to avoid overfitting.

Creating Class Prototypes We used ARCH to extract prototypes of mitotic figures by keyword search. The “with” keyword set included {arrow (arrowhead, circle), mitotic (mitoses)}, as mitotic figures are often highlighted with arrows or circles (see examples in Fig. 3 (a-b)). In total 19 images were identified, from




Fig. 2: (a-b) Results on CAMELYON16 dataset: mIoU (Tumor) and annotated tumor area (%) as functions of annotated tissue area (%). (c) Results on MITOS_WSI_CMC dataset: F1 and the ratio of annotated mitotic figures as functions of annotated tissue area (%). Prototype (adapt) can have different amounts of annotated area as the size of each selected region is dynamically determined. All other methods select regions of size $l \times l$. Random sampling can lead to a slightly smaller annotated area when no more non-overlapping region containing at least 10% tissue is found. All results show median values from five repetitions.

each a section including the mitotic figure was manually cropped. We resized all sections to 64×64 pixels and calculated features using the ResNet18 backbone.

WSI Similarity Map Calculation Patches of size 64×64 pixels ($0.25 \frac{\mu\text{m}}{\text{px}}$, a stride of 64 pixels) were extracted to include only one or a few cells per patch.

Comparison Methods The same settings in the CAMELYON16 experiments were used, using the ResNet18 backbone to extract region features.

Evaluation Scenarios and Metrics The same settings in the CAMELYON16 experiments were used. The three region sizes $l \in \{4096, 8192, 12288\}$ pixels correspond to annotation areas of $\{1.05, 4.19, 9.44\} \text{ mm}^2$, enclosing the recommended area for performing mitotic count (10 high power field (HPF), 2.37 mm^2). The F1 score was used as the evaluation metric, following [2].

3.4 Results

Figure 2 shows the results across all methods. Given a certain budget for annotation area, prototype sampling is more effective than random and diversity sampling in creating a labeled set that maximizes model outcome on both tasks of breast cancer metastases segmentation and mitotic figure detection. This could be attributed to the increased ratios of tumor area and mitotic figures identified




Fig. 3: (a-b) Two example prototype images of mitotic figure. (c-d) An example WSI and its similarity map. Ground truth mitotic figures marked green in (c) and red in (d). The blue boxes in (c) and (d) indicate regions selected with random and prototype sampling, respectively (hyperparameter: 3_M).

and annotated. Notably, prototype sampling achieves near full annotation performance in CAMELYON16 (ViT_PLIP) and MITOS_WSI_CMS (ResNet18) experiments by annotating less than 20% of the tissue area, without the iterative process of annotation region identification in the conventional AL procedures, leading to substantially reduced computational costs. Specifically, on the CAMELYON16 dataset, prototype sampling shows a smaller performance benefit over the other two methods when selecting a large number of regions per WSI, (e.g., in the cases of 3_L and 5_M). This might be because multiple similar regions all possessing tumor characteristics are selected from the same WSI, leading to low data diversity. A comparison between the adaptive and standard region selection methods in Fig. S1 provides further evidence to support this hypothesis: Annotation regions selected using “prototype (standard)” contain a larger amount of tumor area than the annotation regions selected by “prototype (adapt)”, but the trained model shows inferior performance in segmenting tumors on the test set. The standard method can identify a larger area of tumor for annotation as it selects regions that have the highest number of patches resembling tumor, while the adaptive method selects annotation regions centered at the patches that are most similar to tumor prototypes. On the MITOS_WSI_CMC dataset, prototype sampling identifies annotation regions containing more mitotic figures than the other two methods, leading to a consistent increase in model performance. Here, “prototype (standard)” also outperforms “prototype (adapt)” as regions with a larger number of patches resembling mitotic figure are identified for annotation (Fig. S1). Figure 3 (d) provides an example of similarity map, where areas with patches that closely resemble mitotic figure prototypes align well with the areas containing ground truth mitotic figure annotations.

4 Discussion and Conclusion

We presented a novel method for selecting valuable annotation regions that allows to effectively train the model without annotating the entire WSIs. These regions are detected as possessing typical characteristics of each task-specific

class based on their strong resemblance to prototypes recognized from image-caption databases by keyword search or text-to-image retrieval. The efficacy of our method has been proven in different tasks and evaluation settings.

Experiments on the breast cancer metastases segmentation task using both ARCH and OpenPath databases allow evaluating the impact of image-caption database quality on method robustness. While ARCH is well-curated with manually selected samples and cleaned captions, OpenPath is only subjected to image quality sampling testing and automatic text cleaning. We expected that prototypes obtained from OpenPath may have more noise, but we see no drawback of using OpenPath in our experiments. This evaluation was prioritized over an examination involving artificially introduced noise patterns (e.g., swapping captions of two randomly selected pairs), as it depicts “real-world” database quality fluctuations. Future work will investigate the impact of prototype set size.

Particularly, our method can aid the annotation of minority classes. While these classes may also have limited presence in image-caption databases, our method help identify rare but existing samples or discover prototypes of related diseases using prior knowledge. Concerns about finding prototypes for rare diseases are further alleviated by ongoing efforts in creating larger and more diverse histopathology image-caption databases to build powerful all-purpose models.

Acknowledgments. We acknowledge support by d.hip campus - Bavarian aim (J.Q. and K.B.), the German Research Foundation (DFG) project 460333672 CRC1540 EBM, project 405969122 FOR2886 Pandora, as well as the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). The hardware is funded by DFG.

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

References

1. Aubreville, M., Bertram, C., Marzahl, C., Gurtner, C., Dettwiler, M., Schmidt, A., et al.: Deep learning algorithms out-perform veterinary pathologists in detecting the mitotically most active tumor region. *Scientific Reports* **10**(1), 16447 (2020)
2. Aubreville, M., Bertram, C.A., Donovan, T.A., Marzahl, C., Maier, A., Klopffleisch, R.: A completely annotated whole slide image dataset of canine breast cancer to aid human breast cancer research. *Scientific data* **7**(1), 417 (2020)
3. Aubreville, M., Stathonikos, N., Bertram, C.A., Klopffleisch, R., Ter Hoeve, N., Ciompi, F., Wilm, F., Marzahl, C., Donovan, T.A., Maier, A., et al.: Mitosis domain generalization in histopathology images—the MIDOG challenge. *Medical Image Analysis* **84**, 102699 (2023)
4. Bejnordi, B.E., Veta, M., Van Diest, P.J., Van Ginneken, B., Karssemeijer, N., Litjens, G., Van Der Laak, J.A., Hermsen, M., Manson, Q.F., Balkenhol, M., et al.: Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama* **318**(22), 2199–2210 (2017)

5. Bertram, C.A., Aubreville, M., Marzahl, C., Maier, A., Klopffleisch, R.: A large-scale dataset for mitotic figure assessment on whole slide images of canine cutaneous mast cell tumor. *Scientific data* **6**(1), 274 (2019)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International Conference on Machine Learning*. pp. 1597–1607. PMLR (2020)
7. Ciga, O., Xu, T., Martel, A.L.: Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications* **7**, 100198 (2022)
8. CW, E.: The value of histological grade in breast cancer: experience from a large study with long-term follow-up. *Histopathology* **19**, 403–410 (1991)
9. Gamper, J., Rajpoot, N.: Multiple instance captioning: Learning representations from histopathology textbooks and articles. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16549–16559 (2021)
10. Hachohen, G., Dekel, A., Weinshall, D.: Active learning on a budget: Opposite strategies suit high and low budgets. In: *International Conference on Machine Learning*. pp. 8175–8195. PMLR (2022)
11. Huang, Z., Bianchi, F., Yuksekgonul, M., Montine, T., Zou, J.: Leveraging medical twitter to build a visual–language foundation model for pathology ai. *bioRxiv* pp. 2023–03 (2023)
12. Jin, Q., Yuan, M., Qiao, Q., Song, Z.: One-shot active learning for image segmentation via contrastive learning and diversity-based sampling. *Knowledge-Based Systems* **241**, 108278 (2022)
13. Jin, X., An, H., Wang, J., Wen, K., Wu, Z.: Reducing the annotation cost of whole slide histology images using active learning. In: *2021 3rd International Conference on Image Processing and Machine Vision (IPMV)*. pp. 47–52 (2021)
14. Kim, Y.J., Jang, H., Lee, K., Park, S., Min, S.G., Hong, C., Park, J.H., Lee, K., Kim, J., Hong, W., et al.: PAIP 2019: Liver cancer segmentation challenge. *Medical image analysis* **67**, 101854 (2021)
15. Lai, Z., Wang, C., Oliveira, L.C., Dugger, B.N., Cheung, S.C., Chuah, C.N.: Joint semi-supervised and active learning for segmentation of gigapixel pathology images with cost-effective labeling. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 591–600 (2021)
16. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2980–2988 (2017)
17. Litjens, G., Bandi, P., Ehteshami Bejnordi, B., Geessink, O., Balkenhol, M., Bult, P., et al.: 1399 H&E-stained sentinel lymph node sections of breast cancer patients: the CAMELYON dataset. *GigaScience* **7**(6), giy065 (2018), data downloaded from the GigaScience database <http://gigadb.org/dataset/100439>.
18. Lüth, C.T., Bungert, T.J., Klein, L., Jaeger, P.F.: Navigating the pitfalls of active learning evaluation: A systematic framework for meaningful performance assessment. In: *37th Conference on Neural Information Processing Systems* (2023)
19. Mackowiak, R., Lenz, P., Ghorri, O., et al.: Cereals-cost-effective region-based active learning for semantic segmentation. *arXiv preprint arXiv:1810.09726* (2018)
20. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics* **9**(1), 62–66 (1979)
21. Pourahmadi, K., Nooralinejad, P., Pirsavash, H.: A simple baseline for low-budget active learning. *arXiv preprint arXiv:2110.12033* (2021)
22. Qiu, J., Wilm, F., Öttl, M., Schlereth, M., Liu, C., Heimann, T., Aubreville, M., Breininger, K.: Adaptive region selection for active learning in whole slide image

- semantic segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 90–100. Springer (2023)
23. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
 24. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4510–4520 (2018)
 25. Shin, G., Xie, W., Albanie, S.: All you need are a few pixels: semantic segmentation with PixelPick. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1687–1697 (2021)
 26. Veta, M., Heng, Y.J., Stathonikos, N., Bejnordi, B.E., Beca, F., Wollmann, T., Rohr, K., Shah, M.A., Wang, D., Rousson, M., et al.: Predicting breast tumor proliferation from whole-slide images: the TUPAC16 challenge. *Medical image analysis* **54**, 111–121 (2019)
 27. Yang, L., Zhang, Y., Chen, J., Zhang, S., Chen, D.Z.: Suggestive annotation: A deep active learning framework for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 399–407. Springer (2017)
 28. Zheng, H., Yang, L., Chen, J., Han, J., Zhang, Y., Liang, P., Zhao, Z., Wang, C., Chen, D.Z.: Biomedical image segmentation via representative annotation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 5901–5908 (2019)