# Deep Instance-Level Hard Negative Mining Model for Histopathology Images

Meng Li[*], Lin Wu[†], Arnold Wiliem, Kun Zhao, Teng Zhang, and Brian Lovell

The University of Queensland, School of ITEE, QLD 4072, Australia
June 25, 2019

**Abstract.** Histopathology image analysis can be considered as a Multiple instance learning (MIL) problem, where the whole slide histopathology image (WSI) is regarded as a bag of instances (*i.e.,* patches) and the task is to predict a single class label to the WSI. However, in many real-life applications such as computational pathology, discovering the key instances that trigger the bag label is of great interest because it provides reasons for the decision made by the system. In this paper, we propose a deep convolutional neural network (CNN) model that addresses the primary task of a bag classification on a histopathology image and also learns to identify the response of each instance to provide interpretable results to the final prediction. We incorporate the attention mechanism into the proposed model to operate the transformation of instances and learn attention weights to allow us to find key patches. To perform a balanced training, we introduce adaptive weighing in each training bag to explicitly adjust the weight distribution in order to concentrate more on the contribution of hard samples. Based on the learned attention weights, we further develop a solution to boost the classification performance by generating the bags with hard negative instances. We conduct extensive experiments on colon and breast cancer histopathology data and show that our framework achieves state-of-the-art performance.

## 1 Introduction

Deep learning has become increasingly popular in the medical imaging area. However, due to high computational cost, working on whole slide histopathology images (WSIs) with gigapixel resolution is challenging. As a consequence, some approaches attempt to divide WSIs into small patches [5], and predict the final diagnosis of a WSI without providing the pixel-level annotations. To alleviate the annotation efforts, multiple instance learning (MIL) is introduced and explored for weakly annotated WSIs [21], where a WSI is considered as a bag, and the patches of that WSI are regarded as instances. Hence, the problem is cast as dealing with a bag of instances for which a single class label is assigned to the WSI. However, providing insights into the contribution of each instance to

---

[*] meng.li@uq.edu.au; lin.wu@uq.edu.au; arnold.wiliem@ieee.org; {k.zhao1,patrick.zhang}@uq.edu.au; lovell@itee.uq.edu.au
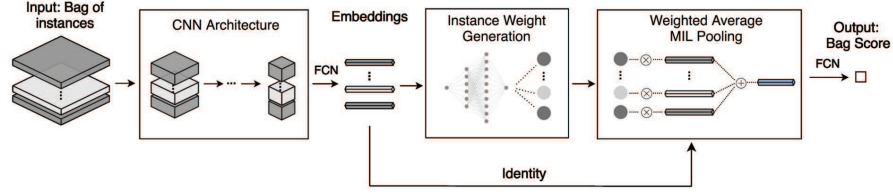[†] Corresponding author

**Fig. 1.** The architecture of the end-to-end deep CNN model with adaptive attention mechanism. The input is the bag of instances (patches of each WSI), which are fed into a CNN model to produce the latent representation of each instance. Then the embeddings of instances go through a fully connected network with the attention weights generated. The learned weights are multiplied with the embeddings of instances in element-wise to be classified by the classifier.

the bag label is crucial to medical diagnosis. Furthermore, if the features of the negative instances are not well learned, the model will be prone to predict a bag with many negative instances to be a positive bag because of the bias issue [4]. In Liu et al.'s work [11], the concept of key instances is discussed, which indicates that in the MIL task, discovering the key instances that trigger the bag label is critical. Thus, one reason for a model to make a wrong prediction about a negative bag is that some challenging instances mislead the model. This will affect decision making of a medical diagnosis model, as a single challenging instance in a WSI may lead to a wrong diagnosis result. The challenging instances herein are referred to as **hard negative patches (instances)** which are incorrectly classified as positive by the prediction model. The other prominent challenge in training a deep model for WSIs is the training imbalance, which means the model tends to become biased towards learning from limited negative patches together with many positive patches.

To address these challenges, we propose an approach that generates the hard negative instances in histopathology images based on instance-level adaptive weighting. Specifically, we incorporate the attention mechanism into the training model, which is computed with adaptive weight loss to explicitly assign larger weights to hard negative instances. The selected instances are used to constitute hard training bags through a bag generation algorithm to further improve the training accuracy. The overview of our framework is shown in Fig. 1. The major contributions of this paper are summarised below.

- We propose a MIL model that incorporates adaptive attention mechanism into an end-to-end deep CNN to detect key instances for histopathology image analysis.
- We present a strategy that is able to address the training imbalance issue through adaptive weighting distribution on negative samples.
- We develop a bag generation algorithm based on the selected hard negative instances and compose the hard bags for the improved classification training.
- Extensive experiments are conducted on real datasets. Experimental results show that our method improves the accuracy, significantly minimizes false positive rate, and also achieves state-of-the-art performance.

## 2   Related Work

**MIL on Histopathology Images** Different MIL approaches were proposed to work on medical images [5,7,10,2]. A two-stage Expectation Maximization based algorithm combined with a deep convolutional neural network (CNN) works well to classify instances on multiple medical datasets [5]. In Kandemir et al.'s work [7], a Gaussian process with relational learning is introduced to exploit the similarity between instances of Barretts cancer dataset. To relate the instances and the bags, different permutation-invariant pooling approaches with CNN have been proposed. One solution is noisy-and pooling function with CNN, which has achieved promising results on medical images [10]. Heather D. Couture et al. [2] adopt an instance-based approach and aggregate the predictions of patches by using a quantile function on breast cancer. In this paper, our work aims to augment the CNN training with attention mechanism to attentively select hard samples, which achieves state-of-the-art performance on colon cancer and breast cancer dataset [16,7].

**Instance-Level Weights** The attention mechanism has become well known in deep learning, especially for natural language processing tasks and image representation learning [9,8,24,23,22]. For MIL problems, attention is represented by the weights of each instance in a bag, a higher weight means more attention. The method to obtain instance weights was first proposed in Nikolaos Pappas and Andrei Popescu-Beliss's work [13], where the weights of instances were trained in a linear regression model. They improved the previous method by using a single layer neural network [14]. A more recent solution was introduced, where the weights were learned by a two-layer neural network [6]. However, these works did not further analyze the generated attention weights. In this work, we propose a method that exploits the relation among the weights and reuses them to improve the model performance.

**Hard Negative Mining** The main idea of hard negative mining is to repeatedly bootstrap negative examples by selecting false positives which the detector incorrectly classifies [3]. Hard negative mining is originally used in object detection tasks, where the datasets usually involve overwhelming easy examples [15]. Recently, this simple but effective technique is commonly used in the medical domain. Difficult negative regions are extracted from the training set to boost model performance on lymph nodes [19] and breast cancer WSIs [1]. Whilst these works make advances in this domain, to our knowledge, there is no work which addresses hard negative mining for MIL tasks. In this paper, we propose an approach which is able to detect hard negative instances by utilizing the attention weight given for each instance in a bag. Based on this, we propose several hard negative bag generation algorithms. As shown in the experiment, our proposed approach achieves significant improvement in comparison with the baselines.

## 3 Our Method

Given a histology image (bag) $\mathbf{X_i}$ with its label $C \in \{0, 1\}$, and its instances $X_{ij} \in \mathbf{X_i}$, $j = 1, 2, ..., N_i$, our goal is to identify a set of hard negative instances from negative bags $\mathbf{X_i}$ with $C_i = 0$. We define $\hat{\mathbf{X}}_\mathbf{i} = \{\mathbf{h_1}, \mathbf{h_2}, ..., \mathbf{h_M}\}$ as the set of hard negative instances identified from the negative bags. We generate hard negative bags from $\hat{\mathbf{X}}_\mathbf{i}$ and use these bags in the training step to improve accuracy. To this end, we develop an instance-level, attention-based CNN model which determines the instance weights by inspecting their corresponding response to contribute the final prediction outcome. The attention selection of key instances allows us to detect the hard negative instances leading to false positives. Then, we propose a bag generation algorithm which produces new bags with selected hard negative patches for retraining.

**Instance-Level Adaptive Weighing Attention** We implement a neural network $f_\phi(\cdot)$ to transform the $j$-th instance in the $i$-th bag into a low-dimensional embedding $g_{ij} = f_\phi(X_{ij})$, where $g_{ij} \in \mathbb{R}^M$. To get the bag representation which is permutation-invariant, a max or a mean aggregation is usually used in the MIL problem [10]. However, neither of them can be trained to get the instance weights which are needed in our case. As a result, a weighted average aggregation method is used. Given the embedding-based instances $G_i = \{g_{i1}, g_{i2}, ..., g_{iN_i}\}$, the weighted average is shown as follows:

$$z_i = \sum_{j=1}^{N_i} w_{ij} g_{ij}, \quad w_{ij} = \frac{\exp(\mathbf{v}_i^\top \frac{\mathbf{U}_i g_{ij}^\top}{1 + |\mathbf{U}_i g_{ij}^\top|})}{\sum_{j=1}^{N_i} \exp(\mathbf{v}_i^\top \frac{\mathbf{U}_i g_{ij}^\top}{1 + |\mathbf{U}_i g_{ij}^\top|})} \tag{1}$$

In (1), $\mathbf{v}_i \in \mathbb{R}^{L \times 1}$ and $\mathbf{U}_i \in \mathbb{R}^{L \times M}$ are the attention network weight parameters. A softsign activation function is implemented to ensure proper gradient flow. We normalize the weights $w_{ij}$ through a softmax function so that all the weights sum to 1. However, the weights defined in (1) are not constrained and could make all the instances receive uniformly distributed weights, which would make it difficult to identify hard negative instances. To this end, we introduce an adaptive weighting method to explicitly enlarge the weights difference between the positive and negative instances:

$$z_i = \sum_{j=1}^{N_i - N_{in}} w_{ij} g_{ij} + \lambda \sum_{j=1}^{N_{in}} w_{ij} g_{ij}, \tag{2}$$

where $N_{in} = |\{X_{ij} | j \in pseudo\ negative\ instances\}|$ is the number of pseudo negative instances. In each iteration, the instances are treated as pseudo negative by thresholding on the average weights in each cycle. The balancing hyperparameter $\lambda$ is empirically set to 2 in our paper. We note that when $\lambda$ is set $> 1$, the weights of the pseudo negative instances will be reduced. This will increase the differences between the positive instance and the negative instances. After obtaining the weighted embeddings, we use a CNN layer for classification. In the
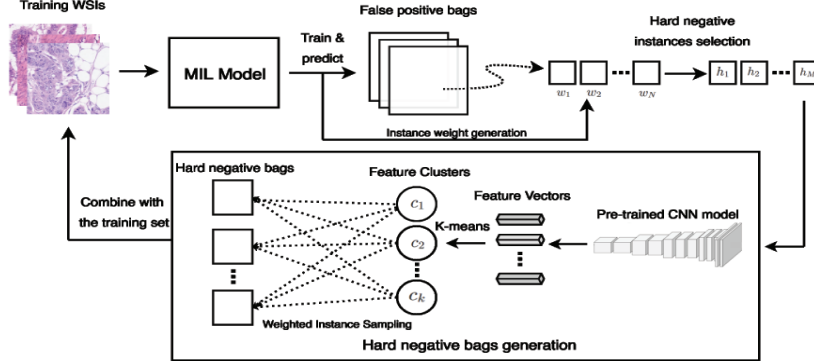
**Fig. 2.** The proposed novel hard negative mining process. The training images are first fed into the deep MIL model with balanced training to select instances to constitute the false positive bags. We learn attention weights for instances, which can be used to select the hard instances that fool the model to make the wrong prediction. Next, the hard negative instances are grouped to form the new hard negative bags by a bag generation algorithm. The patches firstly go through a pre-trained CNN model without the last layer to produce the feature vectors. A K-means clustering is then used to group features into clusters from which hard negative bags are produced by randomly sampling instances across feature clusters dynamically. Finally, the training is augmented with hard negative bags for improved accuracy.

end, the weights of instances are learned so that we can make use of them to mine the hard negative instances.

**Hard Negative Instance Mining** After the training process finishes, we obtain the false positive bags $\{B_1, B_2, ..., B_N\}$, where each bag $B_l$ has the corresponding weights of instances as $\{w_{l1}, w_{l2}, ..., w_{lN_l}\}$, via selecting the hard negative instances through attention weights:

$$H_l = \{w_{li} \mid w_{li} \geq \sigma_l + \overline{w_l}\} \tag{3}$$

where $\sigma_l$ is the standard deviation and $\overline{w_l}$ is the mean of the weights in the $l$-th bag. We group all hard negative instances from all false positive bags together and obtain the bag $\hat{X}_i = \{\mathbf{h_1}, \mathbf{h_2}, ..., \mathbf{h_M}\}$.

**Hard Negative Bag Generation** In this stage, new hard negative bags are generated for re-training the network to further improve the bag prediction accuracy. As histology images contain a variety of patterns in negative regions [20], we propose to select diverse patches into each negative bag such that the deep model will learn features comprehensively. The overview of the bag generation process is shown in Fig. 2. The hard negative instances first go through a pre-trained CNN model to convert to feature vectors. The features are then clustered by the K-means algorithm into $c$ feature groups $C_1, C_2, ..., C_c$, which contain $N_1, N_2, ..., N_c$

items. To conform to the training set bag size, we empirically generate the $i$-th new bag $B_i$ which has a Gaussian random size with $\sigma$ and $\mu$, where $\sigma$ is the standard deviation and $\mu$ is the mean of all training bag sizes. We limit the size from $Z_{min}$ to $Z_{max}$ where $Z_{min}$ and $Z_{max}$ are the minimum and maximum bag size respectively of the training set. We follow a weighted sampling strategy that randomly selects instances from each cluster and puts in $B_i$. The possibility to designate a cluster to sample instance is denoted as $P_j = \frac{N_j}{\sum_{j=1}^{c} N_j}$.

## 4   Experiments

**Datasets** We conduct experiments on two public datasets: Colon Cancer dataset [16] and UCSB dataset [7]. Colon Cancer dataset involves 100 Hematoxylin and Eosin stained (H&E) images from 9 patients at $0.55\mu$m/pixel resolution. The histology images include multiple tissue appearance that belongs to both normal and malignant regions, and can be either used for detection or classification. In total, 22,444 nuclei are annotated in four classes, *i.e.* epithelial, inflammatory, fibroblast, and miscellaneous. A histology image is positive if it consists of at least one epithelial nucleus. We divide each image into $27 \times 27$ patches. The UCSB dataset contains 58 H&E stained image excerpts (26 malignant, 32 benign) from breast cancer patients. The histology image size is $896 \times 768$ and we divide each image into $32 \times 32$ patches. To reduce the noisy patches, each image is converted from the RGB color space to the HSV space, and an Otsu algorithm is used to select the optimal threshold values in each channel to filter the patches [12]. The patches are randomly flipped and rotated for augmentation. Finally, the color normalization (histogram equalization) is performed on each patch.

**Implementation Details** We adopt a deep CNN architecture [16] and a MIL pooling layer [6] to extract features and a fully connected layer to make the classification. More detailed treatment for the architecture can be found in the supplementary material. We follow the work [6], and train the network using an Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. A cross-entropy loss function is used for regression. For the Colon Cancer/UCSB dataset, we set the learning rate $= 5\times10^{-5}/5\times10^{-6}$ and weight decay $= 5\times10^{-4}/1\times10^{-4}$. Each experiment runs for 120/300 epochs, and the epoch with the lowest loss is chosen for evaluation. We evaluate our method using 10-fold/4-fold cross-validation and five repetitions for each experiment. To ensure fair comparison, we test on the same set of bags in each fold for the methods.

**Results and Discussion** Table 1 and Table 2 show the results of our method against different baselines on two datasets. We compare different hard negative mining strategies, *i.e.*, the single hard negative bag generation (SB), the randomly generated multiple bags (MB), and the features clustering bags (FMB). Among these methods, the proposed FMB solution has an overall better result

**Table 1.** Evaluation of different methods for colon cancer data. The average of five experiments with its corresponding standard error is reported. SB: single bag hard negative mining, MB: multiple bags with random sampling, FMB: multiple bags with features clustering sampling, AUC: area under the curve, FPR: false positive rate.

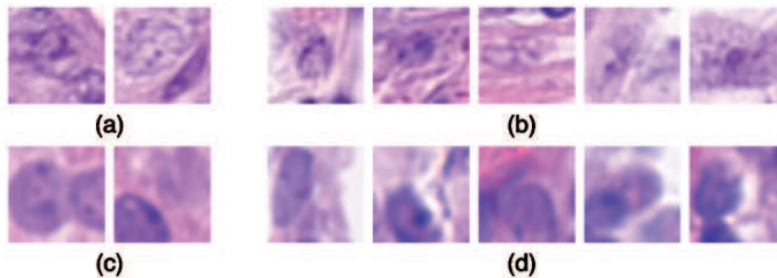| Method | Accuracy | Precision | Recall | F-score | AUC | FPR |
|---|---|---|---|---|---|---|
| MIL Model [6] | 0.904±0.011 | 0.953±0.014 | 0.855±0.017 | 0.901±0.011 | 0.968±0.009 | NA |
| Our Model | 0.906±0.007 | 0.912±0.010 | 0.916±0.012 | 0.905±0.008 | 0.952±0.012 | 0.104±0.012 |
| Our Model+SB | 0.922±0.004 | 0.937±0.011 | 0.920±0.014 | 0.920±0.005 | 0.979±0.004 | 0.084±0.012 |
| Our Model+MB | 0.942±0.005 | 0.963±0.010 | **0.928**±0.017 | 0.939±0.006 | 0.982±0.009 | 0.052±0.010 |
| Our Model+FMB | **0.948**±0.004 | **0.980**±0.003 | **0.920**±0.006 | **0.945**±0.004 | **0.983**±0.004 | **0.036**±0.007 |



**Fig. 3.** Positive and hard negative examples: (a) Colon data: instances that include malignant regions. (b) Colon data: detected hard negative instances that mislead the model to predict a normal bag into a malignant result. (c) UCSB data: instances that include malignant regions. (d) UCSB data: detected hard negative instances that mislead the model to predict a normal bag into a malignant result.

than the others and sets the new state-of-the-art result for colon cancer dataset. For the UCSB dataset, We achieve the same area under the curve (AUC) result as the state-of-the-art method in Song et al.'s work [17]. Fig. 3 demonstrates the detected hard negative instances by our solution. It is noticeable that compared to the baseline method, the false positive rate (FPR) is significantly decreased, and the recall is increased by including hard negative mining, which is particularly essential in histopathology image diagnosis because both false positive and false negative could result in severe consequences for patients. It is also critical that the AUC is increased, as this indicates that the generated hard negative bags do not cause class imbalance issues.

## 5 Conclusions

In this paper, we introduce an effective approach to MIL tasks on histopathology data that incorporates the attention mechanism with adaptive weighing into the deep CNNs for balanced training. To further improve the training accuracy, we develop a novel hard negative mining strategy by generating the bags with hard negative instances. Experimental results demonstrate that our approach makes a decent improvement from the baseline method and achieves state-of-the-art

**Table 2.** Evaluation of different methods for the UCSB dataset. The average of five experiments with its corresponding standard error is reported. SB: single bag hard negative mining, MB: multiple bags with random sampling, FMB: multiple bags with features clustering sampling, AUC: area under the curve, FPR: false positive rate.

| Method | Accuracy | Precision | Recall | F-score | AUC | FPR |
|---|---|---|---|---|---|---|
| MIL Model [6] | 0.755±0.016 | 0.728±0.016 | 0.731±0.042 | 0.725±0.023 | 0.799±0.020 | NA |
| SDR+SVM [17] | **0.983** | NA | NA | NA | **0.999** | NA |
| Our Model | 0.821±0.006 | 0.870±0.032 | 0.800±0.038 | 0.806±0.011 | 0.942±0.008 | 0.220±0.058 |
| Our Model+SB | 0.893±0.001 | 0.920±0.012 | 0.886±0.024 | 0.887±0.015 | 0.949±0.006 | 0.140±0.019 |
| Our Model+MB | 0.936±0.004 | 0.936±0.011 | 0.943±0.014 | 0.935±0.005 | 0.967±0.004 | 0.100±0.016 |
| Our Model+FMB | **0.975**±0.008 | **0.975**±0.012 | **0.979**±0.009 | **0.975**±0.007 | **0.999**±0.001 | **0.040**±0.019 |

performance. In future work, we plan to evaluate our model on the larger WSIs. Moreover, we intend to generalize our model to the multi-class task not limited in the binary classification.

# References

1. Babak Ehteshami Bejnordi, Jimmy Lin, Ben Glass, Maeve Mullooly, Gretchen L Gierach, Mark E Sherman, Nico Karssemeijer, Jeroen Van Der Laak, and Andrew H Beck. Deep learning-based assessment of tumor-associated stroma for diagnosing breast cancer in histopathology images. In *ISBI*, 2017.
2. Heather D Couture, James Stephen Marron, Charles M Perou, Melissa A Troester, and Marc Niethammer. Multiple instance learning for heterogeneous images: Training a cnn for histopathology. In *MICCAI*, 2018.
3. Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
4. Yanjun Han, Qing Tao, and Jue Wang. Avoiding false positive in multi-instance learning. In *NIPS*, 2010.
5. Le Hou, Dimitris Samaras, Tahsin M. Kurc, Yi Gao, James E. Davis, and Joel H. Saltz. Patch-based Convolutional Neural Network for Whole Slide Tissue Image Classification. In *CVPR*, 2016.
6. Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *ICML*, 2018.
7. Melih Kandemir, Chong Zhang, and Fred A. Hamprecht. Empowering multiple instance histopathology cancer diagnosis by cell graphs. In *MICCAI* 2014.
8. Lin Wu, Yang Wang, Junbin Gao and Xue Li. Where-and-When to Look: Deep Siamese Attention Networks for Video-based Person Re-identification. In *IEEE Transactions on Multimedia*, 2019.
9. Lin Wu, Yang Wang, Xue Li and Junbin Gao. Deep Attention-based Spatially Recursive Networks for Fine-Grained Visual Recognition. In *IEEE Transactions on Cybernetics*, 49(5):1791–1802, 2019.
10. Oren Z Kraus, Jimmy Lei Ba, and Brendan J Frey. Classifying and segmenting microscopy images with deep multiple instance learning. In *Bioinformatics*, 2016.

11. Guoqing Liu, Jianxin Wu, and Zhi-Hua Zhou. Key instance detection in multi-instance learning. In *ACML*, 2012

12. Nobuyuki Otsu. A threshold selection method from gray-level histograms. In *SMCS*, 1979.

13. Nikolaos Pappas and Andrei Popescu-Belis. Explaining the stars: Weighted multiple-instance learning for aspect-based sentiment analysis. In *EMNLP*, 2014.

14. Nikolaos Pappas and Andrei Popescu-Belis. Explicit document modeling through weighted multiple-instance learning. In *JAIR*, 2017.

15. Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.

16. K. Sirinukunwattana, S. E. A. Raza, Y. Tsang, D. R. J. Snead, I. A. Cree, and N. M. Rajpoot. Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. In *T-MI*, 2016.

17. Y. Song, Q. Li, H. Huang, D. Feng, M. Chen, and W. Cai. Low dimensional representation of fisher vectors for microscopy image classification. In *T-MI*, 2017.

18. M. Sun, T. X. Han, , and A. Khodayari-Rostamabad. Multiple instance learning convolutional neural networks for object recognition. In *ICPR*, 2016.

19. Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H Beck. Deep learning for identifying metastatic breast cancer. *arXiv preprint arXiv:1606.05718*, 2016.

20. Yan. Xu, J. Zhu, E. Chang, and Z. Tu. Multiple clustered instance learning for histopathology cancer image classification, segmentation and clustering. In *CVPR*, June 2012.

21. Yan. Xu, Tao Mo, Qiwei Feng, Peilin Zhong, Maode Lai, I Eric, and Chao Chang. Deep learning of feature representation with multiple instance learning for medical image analysis. In *ICASSP*, 2014.

22. Lin Wu, Yang Wang, Junbin Gao, and Xue Li. Deep adaptive feature embedding with local sample distributions for person re-identification. *Pattern Recognition*, 73:275–288, 2018.

23. Lin Wu, Yang Wang, Xue Li, and Junbin Gao. Deep attention-based spatially recursive networks for fine-grained visual recognition. *IEEE Transactions on Cybernetics*, 49(5):1791–1802, 2019.

24. Lin Wu, Yang Wang, and Ling Shao. Cycle-consistent deep generative hashing for cross-modal retrieval. *IEEE Transactions on Image Processing*, 28(4):1602–1612, 2019.