



Improving Interpretability for Computer-Aided Diagnosis Tools on Whole Slide Imaging with Multiple Instance Learning and Gradient-Based Explanations

Antoine Pirovano^{1,2(✉)}, Hippolyte Heuberger¹, Sylvain Berlemont¹,
Saïd Ladjal², and Isabelle Bloch²

¹ Keen Eye, Paris, France

{antoine.pirovano,hippolyte.heuberger,sylvain.berlemont}@keeneye.tech

² LTCI, Télécom Paris, Institut Polytechnique de Paris, Paris, France
said.ladjal@telecom-paristech.fr, isabelle.bloch@telecom-paris.fr

Abstract. Deep learning methods are widely used for medical applications to assist medical doctors in their daily routines. While performances reach expert's level, interpretability (highlight how and what a trained model learned and why it makes a specific decision) is the next important challenge that deep learning methods need to answer to be fully integrated in the medical field. In this paper, we address the question of interpretability in the context of whole slide images (WSI) classification. We formalize the design of WSI classification architectures and propose a piece-wise interpretability approach, relying on gradient-based methods, feature visualization and multiple instance learning context. We aim at explaining how the decision is made based on tile level scoring, how these tile scores are decided and which features are used and relevant for the task. After training two WSI classification architectures on Camelyon-16 WSI dataset, highlighting discriminative features learned, and validating our approach with pathologists, we propose a novel manner of computing interpretability slide-level heat-maps, based on the extracted features, that improves tile-level classification performances by more than 29% for tile level AUC.

Keywords: Histopathology · WSI classification · Explainability · Interpretability · Heat-maps

1 Introduction

Since their successful application for image classification [1] on ImageNet [2], deep learning methods (especially Convolutional Neural Network (CNN) deep architectures) have been extensively used and adapted to tackle efficiently a wide range of health issues [3, 4].

Along with these new methods, the recent emergence of Whole Slide Imaging (WSI), microscopy slides digitized at a high resolution, represents a real opportunity for the development of efficient Computer-Aided Diagnosis (CAD) tools to assist pathologists in their work. Indeed, over the last three years, notably due to the WSI publicly available datasets, such as Camelyon-16 [5] and TCGA [6], and in spite of the very large size of these images (generally around 10 giga pixels per slide), deep learning architectures for WSI classification have been developed and proved to be really efficient.

In this work, we are interested in WSI classification architectures that use only the global label (e.g. diagnosis) to train and require no intermediate information such as cell labeling or tissue segmentation (which are time-consuming annotations). The training is regularized by introducing prior knowledge by design in the architectures which, in addition, makes the result interpretable. But the interpretability beyond the architectural design is still pretty shallow.

However, interpretability (capacity to provide explanations that are relevant and interpretable by experts in the field) for medical applications are critical in many ways. (i) For routine tools where useful features are well known and are subject to a consensus among experts, it is important to show that the same features are used by the trained model in order to gain confidence of practitioners. (ii) A good explainability would enable to get the most out of the architectural interpretability and thus assist more efficiently medical doctors in their slide reviews. (iii) The ability to train using only slide level supervision opens a new field we call discovery which consists in predicting, based on easier access (e.g. less intrusive) data, outputs that generally requires heavy processes or waiting such as surgery (e.g. prognosis, treatment response). In order to be able to guide experts towards new discoveries the need for reliable interpretability is obviously high.

In this work, after formalizing the architectural design of most WSI architectures, we propose a piece-wise interpretability approach, that provides cell-level features that prove to be highly relevant and interpretable by pathologists. We also propose a new way of computing explanation slide-level heat-maps based on cell-level identified features and measure their interpretability relevance.

2 Related Work and Motivations

All successful WSI classification architectures deal with these very large images by cutting them into tiles, which is close to the workflow of pathologists who generally analyze these slides at levels of magnification between 5X and 40X.

Recently, as explained in Sect. 1, architectures that are able to learn using only global slide-level labels have been proposed. They rely on a context of Multiple Instance Learning (MIL), i.e. slides are represented by bags of tiles with positive bags containing at least one positive tile and negative bags containing only negatives tiles. For example, CHOWDER [7] is an extension of WELDON [8] solution for WSI classification that uses min-max strategy to guide the training and make the decision. This approach reaches an AUC of 0.858 on Camelyon-16 and 0.915 on TCGA-lung (subset of TCGA dataset related to lung cancer).

In [9], an attention module [10] is used instead of a min-max layer. AUC of 0.775 for a breast cancer dataset and 0.968 for a colon cancer dataset were reported. Recently, more works on large datasets proposed architectures that follow the same design [11, 12]. Heat-maps based on intermediate scores computed in these architectures are what we call architectural explainability that results from prior knowledge on WSI problems that is introduced by design in the architecture. They are of great interest and have proved to be really efficient to the point of being able to spot cancerous lesions that had been missed by experts (in [11]). However explanations are relying on a single “medical” score which might limit the interpretability regarding complex tissue structures that can be found on these slides.

While interpretability for deep learning CNN models is still at its beginning, some methods arise from the literature. “Feature Visualization” has been extensively developed in [13]. It consists of methods that aim at outputting visualizations to express in the most interpretable manner features associated with a single neuron or a group of neurons. It can be used to understand the general training of a model. For example, the question of transferring features learned from natural images (ImageNet) to medical images has only recently been investigated [14] while widely used and yet surprisingly good. It has also been used to measure how robust a learned feature is [15]. Other explainability methods are called attribution methods, i.e. methods that output values reflecting, for each input, its contribution to the prediction. They are performed either through perturbation [16] or gradient computation (i.e. measure of the gradient of the output with respect to the input). This second group of methods is gaining more and more interest. In [17], the authors show that gradient is a good approximation of the saliency of a model and even put forward a potential to perform weakly supervised localization. This work opened a new way of accessing explanations in deep neural networks and motivated a lot of interesting researches [18–20]. Mixed together these explanation methods can provide meaningful and complementary interpretability.

To the best of our knowledge, a lot of explainability is still to be introduced in WSI classification architectures. In the next section, we present our approach to improve interpretability of a model trained for WSI classification in histopathology. We rely on gradient-based methods to identify and attribute the importance of features in intermediate descriptors, and on patch visualization for cell-level feature explanations. We also extend feature explanation to a slide level, thus drastically improving tumor localization and medical insights.

3 Proposed Methods

As introduced in Sect. 2, WSI classification architectures have a common design that we formalize here. Let i be the slide index and j the tile index for each slide. There are four distinct blocks in a typical WSI classification architecture:

1. A feature extractor module f_e (typically a CNN architecture) that encodes each tile $x_{i,j}$ into a descriptor $d_{i,j} \in \mathbb{R}^N$ with N the descriptor size (depending on the feature extractor): $d_{i,j} = f_e(x_{i,j})$;

2. A tile scoring module f_s that, based on each tile descriptor $d_{i,j}$, assigns a single score per tile $s_{i,j} \in \mathbb{R}$: $s_{i,j} = f_s(d_{i,j})$;
3. An aggregation module f_a that, based on all tile scores $s_{i,j}$, and sometimes their tile descriptors $d_{i,j}$, computes a slide descriptor $D_i \in \mathbb{R}^M$ with M the slide descriptor size (depending on the aggregation module): $D_i = f_a(s_{i,j}, d_{i,j})$;
4. A decision module f_{cls} that, based on the slide descriptor D_i , makes a class prediction $P_i \in \mathbb{R}^C$ with C the number of classes: $P_i = f_{cls}(D_i)$.

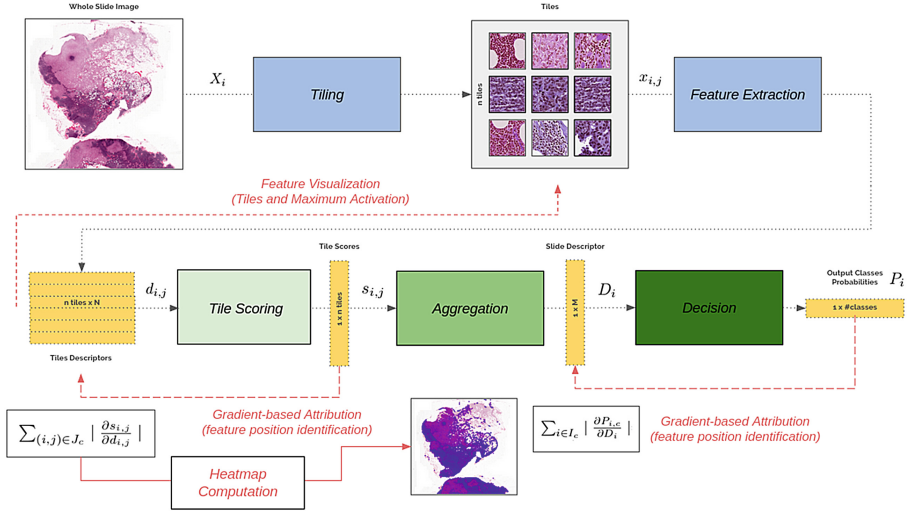


Fig. 1. Overview of the proposed method.

Our approach (illustrated in Fig. 1) consists in rewinding explanations from the decision module to tile information by applying interpretability methods and by answering successively the following three questions:

1. Which features of slide descriptors are relevant for a class prediction?
2. With regards to the aggregation module, which features of tile descriptors are responsible for previously identified relevant slide descriptor features?
3. Are these features of tile descriptors relevant medically and representative of histopathological information?

The first question is answered using attribution vectors $A_c \in \mathbb{R}^M$ (one for each class c) computed as the gradient of the component of index c of P_i (noted $P_{i,c}$) with respect to D_i . It enables us to identify a set of relevant positions $K_c = \{K_{c,1}, \dots, K_{c,L}\}$ in slide descriptors, i.e. the L (empirically determined) positions in A_c with highest attributions over the slide predicted in class c :

$$A_c = \sum_{i \in I_c} \left| \frac{\partial P_{i,c}}{\partial D_i} \right| = \sum_{i \in I_c} \left| \frac{\partial f_{cls}(D_i)_c}{\partial D_i} \right|,$$

with I_c the set of slides predicted to be in class c .

Then, the second question is also answered using an attribution vector $a_c \in \mathbb{R}^N$ computed as the gradient of tile score $s_{i,j}$ with respect to tile descriptor $d_{i,j}$. This enables to identify features positions $k_c = \{k_{c,1}, \dots, k_{c,l}\}$ in tile descriptors, i.e. the l (empirically determined) tile descriptors that are responsible for high activation at previously identified K_c positions in slide descriptor:

$$a_c = \sum_{(i,j) \in J_c} \left| \frac{\partial s_{i,j}}{\partial d_{i,j}} \right| = \sum_{(i,j) \in J_c} \left| \frac{\partial f_s(d_{i,j})}{\partial d_{i,j}} \right|$$

with J_c the set of tile positions (i, j) that most activate K_c positions in slide descriptors.

To answer the third question, we rely on feature activation to highlight features identified as being discriminative to the task by selecting tiles $x_{i,j}$ that have the highest activation per feature in k_c identified over the whole test set. Along with these tiles, we display a maximum activation \mathcal{X} image obtained by iteratively tuning pixels values to activate the feature by gradient ascent as follows: for each k in k_c , \mathcal{X} is initialized as a uniformly distributed image \mathcal{X}_0 ; then while $f_e(\mathcal{X}_{n-1})_k$ increases, iterate over $n > 0$:

$$\mathcal{X}_n(k) = \mathcal{X}_{n-1} + \frac{\partial f_e(\mathcal{X}_{n-1})_k}{\partial \mathcal{X}_{n-1}}$$

Finally, we also propose a new way to compute heat-maps for each slide i . We note $H_{c,i}$ the map that highlights regions on slide i that explain what has been learned to describe class c based on the identified features. For each slide i and tile j , the heat-map value $H_{c,i,j}$ is computed as the average of activations $d_{i,j,k}$ (normalized per feature over all tiles of all slides) over identified features k in k_c for class c :

$$H_{c,i,j} = \frac{1}{|k_c|} \cdot \sum_{k \in k_c} \frac{d_{i,j,k} - \min_k}{\max_k - \min_k}$$

with $\max_k = \max_{i,j}(d_{i,j,k})$ and $\min_k = \min_{i,j}(d_{i,j,k})$.

This heat-map values (between 0 and 1) can be considered as a prediction scoring system, and thus we propose to compute the Area Under the ROC (Receiver Operating Characteristic) Curve to measure how relevant is the interpretability brought by our automatic feature extraction approach using ground truth lesion annotations when given. This localization AUC measures the separability between the class of interest (e.g. ‘‘tumor’’ and other classes using heat-maps, indeed for a good heat-map we expect tiles that are representative of the class of interest to have a high score and all other tiles to have a low score.

4 Experiments and Results

Architectures. We validate our approach on two WSI classification trained architectures: CHOWDER and Attention-based classification.

CHOWDER [7] uses a 1×1 convolution layer to turn each tile descriptor into a single tile score. These scores are then aggregated using a min-max layer, that keeps the top- R and bottom- R scores (e.g. empirically $R = 5$ gives the best results), to give a slide descriptor ($M = 2 \times R$).

Attention-based architecture [9] uses an attention module (two 1×1 convolution layers with respectively 128 and 1 channels and a softmax layer) to compute competitive and normalized (sum to 1) tile scores from tile descriptors. Then, the slide descriptor is computed as the weighted (by tile scores) sum of tile descriptors ($M = N$).

Note that in our experiments the feature extractor is a ResNet-50 [21] ($N = 2048$) trained on ImageNet and the decision module is a two layers fully connected network with 200 and 100 hidden neurons, respectively.

Datasets. We validate our approach using Camelyon-16 dataset that contains 345 WSI divided into 209 “normal” cases and 136 “tumor” cases. This dataset contains slides digitized at 40X magnification from which we extract, with regard to a non-overlapping grid, 224×224 pixels at 20X magnification and pre-compute 2048-tile descriptors for each tile (using the ResNet-50 model trained on ImageNet). 216 slides are used to train our models while 129 slides form the test set to evaluate performances of models.

Results on CHOWDER. Both architectures trained on Camelyon-16 show similar classification performances (AUC of 0.82 for the CHOWDER model and 0.83 for the Attention-based model). Let us now illustrate and detail the results of our approach on the CHOWDER model guided by the three questions raised in Sect. 3.

The first question is “Which slide descriptors features are relevant for a class prediction?” i.e. for CHOWDER given the $M = 10$ ($R = 5$) tile scores given as slide descriptor (the 5 minimum tile scores and the 5 maximum tile scores), what is the contribution of each of these values to the prediction?

The distribution of the (5-)min and (5-)max scores w.r.t. predictions over the whole 129 test slides shows that min scores are the ones that contribute to discriminate between the two classes (i.e. the lower min scores, the more the slide is predicted as being “tumor”). A Mann-Whitney U-Test between scores (min and max independently) distributions reveals that min scores distributions per predicted class are statistically different ($p < 10^{-3}$) while max scores are not ($p = 0.23$). The attribution of min and max scores distributions validates this assertion.

After explaining that min scores are the ones describing tumorous regions and thus that max scores are used for the “normal” class, we are interested in identifying which features of tile descriptors are mostly responsible for minimum and maximum scores, i.e. to describe each class. To address this second question, we use the same gradient-based explanation method.

Most minimal tile scores are under -5 and most maximal tile scores are above 11. For each of these groups of tiles, we compute the average attribution of each of the $N = 2048$ features in tile descriptors (extracted by a ResNet-50 trained on ImageNet). The distribution of features hence activated allows us to identify which features are mostly responsible for min and max tile scores, i.e. highest attribution for min and max scored tiles.

Thus we are able to claim that features (defined by their position in the descriptor) that are mostly useful for the trained model for each class are: 242, 420, 602, 1154, 1644, 1652 and 1866 for “tumor” class, and 565, 628, 647, 1158 and 1247 for “normal” class.

Interpretability. As exposed in the previous paragraph, based on explanations on decision blocks, we have been able to identify 7 and 5 features that are mostly used by the trained CHOWDER model to make decisions (and we did the same for the attention-based model). Now, we are interested in interpretable information to return to pathologists so that they can use their expertise to understand what these features put forward histopathologically speaking. We benefited from discussions with two experienced pathologists and report their overall feedback on the interpretable visualization we proposed.

Figure 2 shows the 7 tiles that activate the most (over all tiles) each feature and the max activation image, that we expect to reveal what the feature reveals with regards to the histopathological problem it has been trained on.




Fig. 2. Patch-based visualizations obtained for features 242, 1154 and 1652 (for min-scores features); 565 and 1247 (for max-scores features); tiles and max activation images (right).

Pathologists agreed that patch-based tiles visualizations are highly interpretable and reveal features that are indeed related to each class [22]. For example, feature 1652 tends to trigger spindle-shaped cells that indeed can be a metas-

tastic tissue organization. For “normal” tissue features, feature 565 describes mainly clustered lymphocytes that are preponderant in normal tissues.

Coherence between patches exposed for a better interpretability led us to think about another way to present features to pathologists. Indeed, since tissues have coherent and somehow organized structure, a relevant feature for histological problems would be activated in a coherent and somehow organized way over slides. Thus, along with patch-based visualization, we propose to access features activation heat-maps $H_{c,i}$ over slides as presented in Sect. 3.

Figure 3 illustrates qualitative results. Quantitatively, we report a tile-level localization AUC of 0.884 for CHOWDER model and 0.739 for Attention-based model, using this average normalized activation as a “tumor” prediction score and using lesion annotation provided by Camelyon-16 dataset to get the ground-truth label per tile. Both AUCs are significantly high, which validates our approach of identifying features that are relevant and of computing heat-maps for interpretation and explanation. Note that the AUC computed using tile scores is 0.684 for CHOWDER model and 0.421 for Attention-based model (see Table 1). We can also note that there is a gap in interpretability between CHOWDER model and Attention-based model while classification performances are comparable. The gap can be explained by the fact that, in the context of Camelyon-16, identifying one tumorous tile is enough to label a slide as “tumor”, so implicit tile classification does not need to be exhaustive to provide meaningful information to the slide level decision module, however if so interpretability will decrease.




Fig. 3. Slide-based visualizations: Heat-maps explaining “tumor” class obtained by computing average normalized activation over identified features; ground-truth annotations for “tumor” tissue (left); CHOWDER model feature-based heat-maps (middle); attention-based model feature-based heat-maps (right).

Table 1. Results: classification and localization AUC using both methods (improvement of localization AUC by 29.2% for CHOWDER and 75.5% for Attention-based model).

Model	Classification AUC	Heat-map method	Localization AUC
CHOWDER	0.82	Tile scores	0.684
		Feature-based (ours)	0.884
Attention-based	0.83	Tile scores	0.421
		Feature-based (ours)	0.739

5 Conclusion

In this paper, we presented our interpretability approach and researches for WSI classification architectures. We proposed a unified design that gathers a large majority of WSI classification methods relying on MIL learning, and applied a gradient-based attribution method to identify features that have been learned to be relevant in intermediate (tile and slide) descriptors. Then we showed the relevance of these features by visualization, and validated it with the help of pathologists. We finally proposed explainability heat-maps over whole slides taking into account only identified features. This contribution considerably improved tile-level classification AUC. Allying patch-based and slide-based visualization took interpretability to a next level for pathologists to understand histological meanings of features used by trained models. More generally, this method could be used to extract meaningful and interpretable information for any medical application when using a MIL context.

Acknowledgment. We would like to gratefully thank Jacques Bosq and Yaelle Harrar for their important medical inputs and feedbacks during this work, along with our colleagues at Keen Eye for their work and support. This work was partially funded by ANRT, Grant N° 2017/1076.

References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, vol. 25, pp. 1097–1105 (2012)
2. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: *Conference on Computer Vision and Pattern Recognition* (2009)
3. Pasa, F., Golkov, V., Cremers, D., Pfeiffer, F.: Efficient deep network architectures for fast chest x-ray tuberculosis screening and visualization. *Nature Sci. Rep.* **9**, 62–68 (2019)
4. Pratt, H., Coenen, F., Broadbent, D.M., Harding, S.P., Zheng, Y.: Convolutional neural networks for diabetic retinopathy. *Procedia Comput. Sci.* **90**, 200–205 (2016)

5. Ehteshami Bejnordi, B., et al.: Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *J. Am. Med. Assoc.* **312**(22), 2199–2210 (2017)
6. Tomczak, K., Czerwińska, P., Wiznerowicz, M.: The cancer genome atlas (TCGA): an immeasurable source of knowledge. *Contemp. Oncol. (Pozn)* **19**(1A), A68–A77 (2015). <https://doi.org/10.5114/wo.2014.47136>
7. Courtiol, P., Tramel, E.W., Sanselme, M., Wainrib, G.: Classification and disease localization in histopathology using only global labels: a weakly-supervised approach. *Computing Research Repository (CoRR)*, Arxiv (2018). <https://arxiv.org/abs/1802.02212>
8. Durand, T., Thome, N., Cord, M.: WELDON: weakly supervised learning of deep convolutional neural networks. In: 29th IEEE Conference on Computer Vision and Pattern Recognition, June 2016
9. Ilse, M., Tomczak, J.M., Welling, M.: Attention-based deep multiple instance learning. In: *Proceedings of the International Conference on Machine Learning (ICML)* (2018)
10. Raffel, C., Ellis, D.P.W.: Feed-forward networks with attention can solve some long-term memory problems. *Computing Research Repository (CoRR)*, Arxiv (2015). <https://arxiv.org/abs/1512.08756>
11. Campanella, G., Hanna, M.G., Geneslaw, L., Mirafior, A., Silva, V.W.K., Busam, K.J.: Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Med.* **25**, 1 (2019). <https://doi.org/10.1038/s41591-019-0508-1>
12. Li, J., Li, W., Gertych, A., Knudsen, B.S., Speier, W., Arnold, C.W.: An attention-based multi-resolution model for prostate whole slide image classification and localization. *Computing Research Repository (CoRR)*, Arxiv (2019). <https://arxiv.org/abs/1905.13208>
13. Olah, C., et al.: The building blocks of interpretability. *Distill* **2**, e10 (2018)
14. Raghu, M., Zhang, C., Kleinberg, J., Bengio, S.: Transfusion: understanding transfer learning for medical imaging. In: *Proceedings of the Neural Information Processing Systems*, pp. 3342–3352 (2019)
15. Couteaux, V., Nempont, O., Pizaine, G., Bloch, I.: Towards interpretability of segmentation networks by analyzing DeepDreams. In: Suzuki, K., et al. (eds.) *ML-CDS/IMIMIC -2019. LNCS*, vol. 11797, pp. 56–63. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-33850-3_7
16. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: *IEEE International Conference on Computer Vision (ICCV)*, pp. 3449–3457, October 2017. <https://doi.org/10.1109/ICCV.2017.371>
17. Simoyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: visualising image classification models and saliency maps. *Computing Research Repository (CoRR)*, Arxiv, December 2013. <https://arxiv.org/abs/1312.6034>
18. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *34th International Conference on Machine Learning*, vol. 70, pp. 3319–3328, March 2017
19. Smilkov, D., Thorat, N., Kim, B., Viégas, F., Wattenberg, M.: SmoothGrad: removing noise by adding noise. *Computing Research Repository (CoRR)*, Arxiv, June 2017. <https://arxiv.org/abs/1706.03825>
20. Goh, G.S.W., Lapuschkin, S., Weber, L., Samek, W., Binder, A.: Understanding integrated gradients with SmoothTaylor for deep neural network attribution. *Computing Research Repository (CoRR)*, Arxiv, April 2020. <https://arxiv.org/abs/2004.10484>

21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778, June 2016. <https://doi.org/10.1109/CVPR.2016.90>
22. Hoon Tan, P., et al., The WHO Classification of Tumours Editorial Board: Breast Tumours, WHO Classification of tumours. International Collaboration on Cancer Reporting (ICCR), 5th edition, December 2019