



An Interactive Approach to Region of Interest Selection in Cytologic Analysis of Uveal Melanoma Based on Unsupervised Clustering

Haomin Chen¹(✉), T. Y. Alvin Liu², Zelia Correa², and Mathias Unberath¹

¹ Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA
hchen135@jhu.edu

² Wilmer Eye Institute, School of Medicine, Johns Hopkins University,
Baltimore, MD, USA

Abstract. Facilitating quantitative analysis of cytology images of fine needle aspirates of uveal melanoma is important to confirm diagnosis and inform management decisions. Extracting high-quality regions of interest (ROIs) from cytology whole slide images is a critical first step. To the best of our knowledge, we describe the first unsupervised clustering-based method for fine needle aspiration cytology (FNAC) that automatically suggests high-quality ROIs. Our method is integrated in a graphical user interface that allows for interactive refinement of ROI suggestions to tailor analysis to any specific specimen. We show that the proposed approach suggests ROIs that are in very good agreement with expert-extracted regions and demonstrate that interactive refinement results in the extraction of more high-quality regions compared to purely algorithmic extraction alone.

Keywords: Human-computer interaction · Unsupervised learning · Machine learning · Coarse to fine

1 Introduction

Clinical Background: Uveal melanoma is the most common primary intraocular malignancy in adults [21]. As standard care for uveal melanoma, Fine Needle Aspiration Biopsy (FNAB) is often performed to confirm the diagnosis and to obtain cell aspirates for both Gene Expression Profile (GEP) and Cytology of Fine Needle Aspirates (FNAC) analysis for prognostication. According to recent analysis, primary uveal melanoma clusters in two distinct subgroups according to its GEP; the first corresponding to low grade melanoma with little to no metastatic risk, and the second corresponding to high grade melanoma with high metastatic risk, which results in 6 times of 5-year probability of metastatic death [7]. While GEP analysis of fine needle aspirates has shown good accuracy for identifying patients at high risk of metastatic disease, the only commercially

available test is expensive, requires special storage and transportation, has a long turn around time and is only available in the US. Most importantly, despite its efficacy, the commercial GEP test still occasionally fails resulting in unpleasant clinical surprises and unexpected early metastatic death. There is increasing evidence that the underlying genetic profile affects cancer growth on multiple scales. Radiomics, for example, exploit this observation to develop imaging-derived biomarkers that are informative for prognosis [12]. We hypothesize that such multi-scale analysis will also be useful for prognosis in uveal melanoma. Specifically in addition to GEP, we would like to extract imaging-features from FNAC. In addition to complementing GEP analysis, such cytology-based test could provide a cheap and widely available alternative for prognostication of uveal melanoma [9]. However, pathologist analysis of FNAC is infeasible, as 1) it is a very time-consuming and tedious task, and 2) none of the manually defined cytopathological features proved particularly robust for predicting metastatic risk.

To reach this goal, we need to facilitate or even automate quantitative analysis of cytology whole slide images (WSIs). To this end, we develop an interactive tool that our envision will be beneficial in two ways: First, it can be deployed in pathologist-centric workflows to guide pathologist review, thereby reducing the experts workload. Second, the tool provides an opportunity for pathologists to guide algorithmic evaluation, e.g. by refining the content that is submitted for automated analysis of the slide, e.g. for GEP classification. Such an interactive design may prove beneficial in building trust, accelerating workflows, and reducing mistakes, of both automated algorithms and pathologists.

We present our first steps in this direction that consider the extraction of high-quality Region of Interests (ROIs) (areas with multiple clear cancer cells) from gigapixel-sized histological architecture, FNAC. We propose a Human-Interactive Computationally-Assisted Tool (HICAT) that supports ROI selection with a 2-step coarse-to-fine unsupervised clustering. It provides **interactivity** to allow for patient-specific refinement of ROI selection at application time. This refinement provides insight in and some control over the region selection, and results in the extraction of more informative regions compared to the purely algorithmic extraction. Such human-machine partnership may contribute to pathologists building trust in AI-assisted tools. HICAT increases Recall in ROIs from 7.44% to 42.32%, while Precision remains the same 83%. On average, 1318 ROIs per FNAC are extracted, which contains enough information for further analysis. Our AI-assisted ROI selection workflow is more than 10 times faster than manual ROI extraction by pathologists that was used previously [17].

Related Work: Histology WSI and FNAC are two main foci in pathology. Histology WSI contains an entire slice of tissue and several learning-based algorithms for ROI extraction have been proposed [3, 14, 15, 19]. FNACs exhibit high variation in cell quality and artifact, and to our knowledge, all existing approaches for high-level FNAC analysis operate on manually identified ROIs [8, 10, 20]. Due to the small targets (e.g. lesions and organs) in medical

imaging, coarse-to-fine concepts are widely used. Spatial coarse-to-fine segmentation is applied to target small organs and lesions [6, 16, 25–27]. Spatial coarse-to-fine clustering is also commonly used to extract ROIs from high spatial resolution WSIs and several machine-learning approaches exist for this task [3, 19]. To involve human cognition and refine algorithmic prediction, human interaction with deep learning so far has been largely limited to segmentation problems [1, 2, 11, 22, 23].

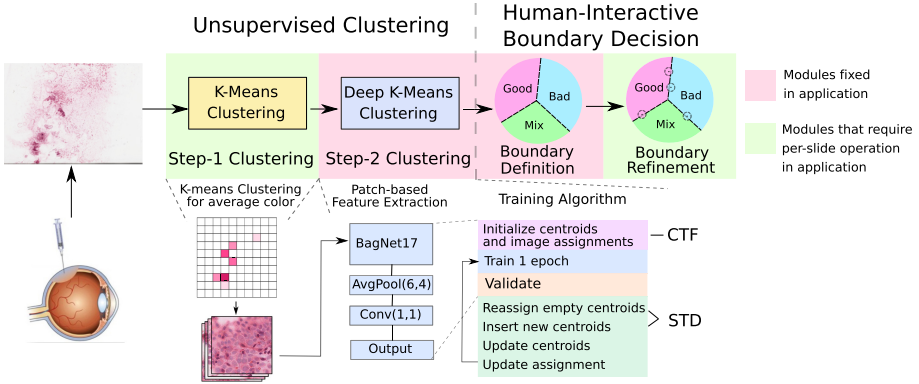


Fig. 1. Overview of the HICAT.

2 Method

Given a FNAC image, we seek to extract square-shaped ROIs, similar to those shown in Fig. 2(a), which lend themselves well for further cell-level algorithmic analysis. Our ROI extraction pipeline contains of a 2-step clustering that is followed by an interactive decision boundary definition to assign image-quality to centroids. The clustering algorithm will be discussed in Sect. 2.1 and Sect. 2.2. The first step aims to remove blank images, *i.e.* Fig. 2(g), to greatly reduce processing time for the second step, which further clusters the selected ROIs based on image content. After the 2-step clustering, a global decision boundary for all FNACs is defined by centroid-level human annotation. Interactive refinement of this decision boundary is then possible for every patient and FNAC to improve the algorithmic ROI selection based on centroid annotation (Sect. 2.3).

2.1 Step-1 Clustering

The given FNAC is first down-sampled such that each pixel in the resulting image corresponds to the average signal within one area. The size of this area is only constrained by its compatibility with the following clustering steps. We

found the size 512×512 is able to perform sufficiently well. K-means clustering is then used to cluster pixel intensities into 2 centroids that intuitively correspond to regions with bright and dark average intensities. Since FNACs are acquired with the bright-field technique, pixels with low and high intensities correspond to regions with high and low tissue content, respectively. We select the darker centroid for further processing via Step-2 clustering in Sect. 2.2. Because the exact magnitude of bright and dark centroid intensities varies with cell distribution and illumination, this scheme is applied to every FNAC slide independently.




Fig. 2. Different types of ROIs in FNACs. (a) High-quality ROIs, which contain more than 3 clear cancer cells. (b) Blood cell ROIs. (c) Blurred ROIs. (d) Fluid ROIs. (e) Multi-layer cell ROIs. (f) Artifact ROIs. (g) Blank ROIs. (h) Borderline ROIs, which contain more than 3 clear cancer cells, but contains a large portion of low-quality areas.




Fig. 3. An example of Step-1 clustered area and some of the corresponding Step-2 clustering ROIs. (a) Step-1 area. (b) Top-left corner. (c) Top-right corner. (d) Bottom-left corner. (e) Bottom-right corner.

2.2 Step-2 Clustering

Step-2 clustering aims to separate high-quality images with more than 3 clear cancer cells from low-quality images that either show blood cells, multiple layers of cells and fluid, are blurred or otherwise corrupted with artifact. Examples of such images are provided in Fig. 2. Since this separation is based on image content that, in cytology, can vary considerably across pixels (cf. Fig. 3), a patch-based network is applied to perform clustering on 228×228 pixel ROIs in naive

resolution which is much smaller than the areas extracted from Step-1 clustering. These patches are extracted with a stride of 128 from the ROIs selected in Step-1 clustering. A previous state-of-the-art patch-based method, BagNet17 [4] is used as the backbone. The input images of size 512×512 pixels are first down-sampled 4 times and an average pooling layer with kernel size 6 and stride 4 is attached after the final residual block, so that each output pixel corresponds to one desired patch (if using other parameters, the receptive field’s size and stride cannot be guaranteed to take on the desired value). Finally, a convolutional layer with kernel size 1×1 compresses the feature into a lower-dimension space with dimension d . We follow [5, 24] to involve k-means clustering for the d -dimension network outputs. K-means centroids and patch assignments are initialized by the pre-trained network and are fixed in the training phase. L2 loss is applied to force patch features to be close to the assigned centroid. Centroids and patch assignments are updated during the validation phase. We reassign empty centroids during training to avoid trivial parametrization. Step-2 clustering is trained on all FNACs simultaneously.




Fig. 4. Examples for FNAC-specific ROI refinement GUI. For each screenshot, top-left image is the down-sampled WSI, top-right image shows the corresponding spatial states for all ROIs, white/light grey/grey means high-/mix-/poor-quality ROIs. Dark grey corresponds to blank images removed by Step-1 clustering. Pink pixels correspond to uncertain ROIs. Bottom left image is the corresponding full resolution ROI that the mouse hovers over. By double clicking the pixel on the down-sampled WSI or state image, a window in bottom right will pop out for annotation. (a) shows the overall behaviour of the state image. (b) shows the zoomed-in version for detail visualization.

In order to reduce the number of centroids that focus on fluid and artifact images, we introduce a centroid-based coarse-to-fine clustering strategy. Only a portion of centroids are initialized first, and new centroids are inserted during training in order to increase the probability of these centroids to account for cell images, which is referred as *CTF* in Fig. 1. We reassign/insert empty/new centroids around the centroid with the largest standard deviation of its assigned samples in feature space, instead of the centroid with the largest number of samples [5, 24]. It is referred as *STD* in Fig. 1. This is because of 2 reasons: 1) A considerable number of fluid and artifact images exists and there is no use to further insert centroids for these images. 2) Fluid and artifact images are easier to separate because of the difference in complexity compared to cell images.

Consequently, centroids with cell images tend to have larger standard deviation among the assigned samples in feature space, so that inserted/re-assigned centroids are more likely to focus on cell images. The re-assignment and insertion is processed during the validation phase.

2.3 Interactive Centroid Assignment and Refinement

After Step-2 clustering, every centroid contains ROIs that exhibit similar appearance. However, at this point it is still unclear which of the ROIs in the centroids are high-/low-quality. To provide this semantic definition with minimal manual annotation requirement, we developed a Graphical User Interface (GUI) that allows for rapid centroid annotation. To this end, 10 ROIs from 10 random centroids are displayed for the user to classify. After several iterations, each centroid has more than 10 high-/poor-quality annotations. The ratio of high-quality ROIs classified to every centroid is then used to define a centroid-level boundary that separates between high- and low-quality ROIs. Because cell quality in FNACs has large variation, some ROIs cannot be clearly classified as high-/low-quality, *e.g.* Fig. 2(h). Therefore, we allow for some mix-quality centroids that contain roughly an equal number of high-/low-quality ROI annotations. Although there exists high-quality ROIs in mix-quality centroids, we exclude them to avoid introducing poor-quality images to influence further analysis.

During application, due to high variations in FNACs, the classifier based on the above procedure may not perform perfectly when suggesting ROIs in new FNACs. To allow for the refinement of ROI suggestions, a patient-specific refinement tool is created for pathologists to interact with, as shown in Fig. 4. Specifically, high-/low-/mix-quality assignments from boundary definition are visualized and synchronized with the corresponding FNAC image. The user can hover the mouse over the FNAC to display the underlying ROI in native resolution, and can simply click it to re-annotate if necessary. In this case, the selected ROI and all ROIs with similar features $\{x, \text{where } \|x - F\|_2 < \lambda_2 L_1\}$ are all re-annotated, where F is the selected ROI's feature, L_1 is the distance to the closest centroid and λ_2 is a constant. Uncertain ROIs are also identified and displayed to users as recommended for re-annotation. Using L_1, L_2 as the distance of an ROI feature to the 2 closest centroids. The ROI is considered uncertain if the two closest centroids are high- and low-quality, respectively, and satisfies

$$\frac{\|L_1 - L_2\|_2}{\min\{L_1, L_2\}} < \lambda_1 \quad (1)$$

where λ_1 is a constant. The result of every click re-annotation is reflected in real time. The user has full control over when to stop the refinement.

3 Experiment

3.1 Experiment Setup

Dataset: The dataset we use includes 100 FNAC samples from 100 uveal melanoma patients. The cellular aspirates obtained from FNACs of each tumor

were submitted to cytology and GEP testing. The cytology specimen was flushed on a standard pathology glass slide, smeared, and stained with hematoxylin and eosin. The specimen submitted for GEP was flushed into a tube containing extraction buffer and submitted for DecisionDx-UM testing. Whole slide scanning was performed for each cytology slide at a magnification of 40x, using the Aperio ScanScope AT machine, and the high-magnification digital image was examined using the Aperio Imagescope software.

516 areas of size 1716×926 are manually extracted and annotated from 20 slides by an expert pathologist. Every area is split into 8 small areas with equal size. Each small area is further split into 9 ROIs where the stride of ROI extraction is half of their width and height. All of these ROIs are annotated as high-/low-quality images, which results in 37, 152 annotated ROIs. The criterion for high-quality images is the same as Fig. 2(a). All our experiments are trained on the remaining 80 slides and tested on the 20 slides with annotations.

Implementation Details: 259, 203 areas are extracted by Step-1 clustering. In Step-2 clustering, each area corresponds to 9 ROIs with size 228×228 , which results in a total of 2, 332, 827 ROIs for training. The length d of the output feature vector is 16. Centroid-based coarse-to-fine clustering is first initialized with 32 centroids. 4 new centroids are inserted after every training epoch until a total of 100 centroids exists. We implement the model using PyTorch [18] for Step-2 clustering, and initialize them with ImageNet pre-trained weights provided by [4]. All models are optimized by Adam [13] with a learning rate of 10^{-3} . All interactive centroid assignments and specific boundary refinement were performed by an expert pathologist. During centroid definition, centroids with greater than 70% of ROIs annotated as high-quality are classified as high-quality centroids, while centroids with fewer than 30% are classified as low-quality centroids. The other centroids are mix-quality centroids. For boundary refinement, the parameters are $\lambda_1 = 0.2, \lambda_2 = 0.5$.

Table 1. Ablation study for clustering algorithm. DeepCluster (DC) is DCN [24] with BagNet17 [4] as backbone. ‘‘CTF’’ indicates the use of the proposed centroid-based coarse-to-fine strategy. ‘‘STD’’ indicates the use of the proposed mechanism of inserting/reassigning new/empty centroids to be around the centroid with the largest standard deviation of its assigned samples in feature space. (Otherwise, to be around the centroid with most samples). Numbers of high-/low-quality centroids are also reported.

Model	Recall _{gb}	Recall _{gmb}	Precision _{gmb}	Accuracy	#high-quality	#low-quality
DC [4, 24]	11.74%	7.44%	83.17%	61.43%	10	60
DC+STD	34.71%	7.89%	85.99%	63.63%	18	43
DC+STD+CTF	51.38%	27.83%	91.56%	70.90%	23	51

Evaluation Metrics: The final goal for our proposed extraction is to maximize the number of high-quality ROIs and to minimize the number of

low-quality ROIs provided for further analysis. To evaluate our success, we calculate the recall, precision and accuracy on the ROIs in the 20 slides with manually extracted ROIs. Because there exist mix-quality centroids, we first report recall and precision for images only in high-/low-quality centroids, denoted as Recall_{gb} , Precision_{gb} . We also report recall, precision and accuracy for all annotated images, by treating mix-quality centroids as low-quality centroids, denoted as Recall_{gmb} , Precision_{gmb} and Accuracy. Because Precision_{gb} is the same as Precision_{gmb} , only Precision_{gmb} is recorded.

3.2 Ablation Study for Clustering Algorithm

In order to compare different clustering algorithms, human-interactive boundary definition is performed separately for all models to classify high-/mix-/low-quality centroids by the same expert pathologist. We conduct an ablation study for clustering algorithm to analyze the contributions of its novel components. The baseline is the combination of the deep clustering network, DCN [24], with Bag-Net17 [4] (referred to as *DeepCluster*) with 100 centroids. The performance by adding the two novel components: centroid-based coarse-to-fine concept (referred to as *CTF*) and the centroid insertion/reassignment algorithm (referred to as *STD*) is compared. The Step-1 clustering is kept the same across all models, which eliminates 96.5% areas as blank areas. Results are summarized in Table 1.

The effect of our proposed centroid insertion/reassignment algorithm is reflected in the comparison of *DeepCluster vs. DeepCluster+STD*. Recall_{gb} and Precision_{gmb} increase from 11.74% and 83.17% to 34.71% and 85.99% by using *STD*. Improvements are due to our observation that standard deviation of the assigned samples are efficient to tell apart centroids for high-/low-quality images. More centroids for high-quality images result in better performance.

The effect of centroid-based coarse-to-fine method is reflected in the comparison of *DeepCluster+STD vs. DeepCluster+STD+CTF*. By adding the centroid-based coarse-to-fine module to *DeepCluster+STD*, we observe substantial improvements in Recall_{gmb} and Precision_{gmb} which increase from 7.89% and 85.99% to 27.83% and 91.56%, respectively. The improvement is in line with our motivation and hypothesis that more centroids are assigned to focus on images with different cells and various visual quality. The increase in the number of high-quality centroids further supports our hypothesis.

Table 2. Ablation study for human interactive patient-specific boundary refinement.

Model	Recall_{gb}	Recall_{gmb}	Precision_{gmb}	Accuracy
Without boundary refinement	51.38%	27.83%	91.56%	70.90%
HICAT	59.47%	42.32%	83.09%	74.18%

3.3 Ablation Study for Interactive Refinement

The performance of interactive refinement of ROI suggestion is shown in Table 2. FNAC and the ROIs' labels after centroid definition is synchronously visualized as Fig. 4. An expert pathologist finished the human interactive boundary refinement for all testing FNACs. Less than 50 re-annotation clicks are performed for each slide. The pathologist stopped the process for each slide, once he determined there were adequate high-quality ROIs selected for further analysis and few low-quality ROIs exist. Comparing with/without boundary refinement shows that Recall_{gmb} goes drastically up from 27.83% to 42.32%. The reduced precision from 91.56% to 83.09% may be attributed to a conservative selection of the pathologist. However, since adequate high-quality ROIs are still available for further analysis, this decrease is likely not problematic. The boost in performance is due to the variation in different FNACs. Pathologists may interact with our tool to adjust the inclusion criteria based on a specific FNAC, *e.g.* when few cells are visible, the selection criteria for high-quality ROIs can be relaxed. Finally, 1318 ROIs are extracted on average per FNAC, which contain adequate information for further analysis. The whole application process takes 15 min per FNAC, which is more than 10 times faster than manual ROI extraction. (3 min for 2-step clustering and 12 min for boundary refinement.)

4 Conclusion

In this paper, we propose an interactive and computationally-assisted tool for high-quality ROI extraction from FNACs. Our method relies on 2-step unsupervised clustering of ROI appearance and content to automatically suggest ROI of acceptable quality. These suggestions can then be refined interactively to adapt ROI selection to specific patients. We hope to contribute effective tools that support quantitative analysis of FNACs to, in the future, improve prognostication of patients suffering from uveal melanoma.

Acknowledgement. We gratefully acknowledge funding from the Emerson Collective Cancer Research Fund and internal funds provided by the Wilmer Eye Institute and the Malone Center for Engineering in Healthcare at Johns Hopkins University.

References

1. Amrehn, M., Gaube, S., Unberath, M., et al.: UI-NET: interactive artificial neural networks for iterative image segmentation based on a user model (2017)
2. Aresta, G., et al.: iW-Net: an automatic and minimalistic interactive lung nodule segmentation deep network. *Sci. Rep.* **9**, 11591 (2019)
3. Barker, J., Hoogi, A., Depeursinge, A., Rubin, D.L.: Automated classification of brain tumor type in whole-slide digital pathology images using local representative tiles. *Med. Image Anal.* **30**, 60–71 (2016)
4. Brendel, W., Bethge, M.: Approximating CNNs with bag-of-local-features models works surprisingly well on ImageNet. In: International Conference on Learning Representations (2019)

5. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. LNCS, vol. 11218, pp. 139–156. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01264-9_9
6. Chang, L., Zhang, M., Li, W.: A coarse-to-fine approach for medical hyperspectral image classification with sparse representation. In: Yu, J., et al. (eds.) *AOPC 2017: Optical Spectroscopy and Imaging*, vol. 10461, pp. 136–144. International Society for Optics and Photonics, SPIE (2017)
7. Corrêa, Z., Augsburger, J.: Sufficiency of FNAB aspirates of posterior uveal melanoma for cytologic versus GEP classification in 159 patients, and relative prognostic significance of these classifications. *Graefe’s archive for clinical and experimental ophthalmology = Albrecht von Graefes Archiv fur klinische und experimentelle Ophthalmologie* 252 (2013)
8. Dov, D., Kovalsky, S.Z., Cohen, J., Range, D.E., Henao, R., Carin, L.: A deep-learning algorithm for thyroid malignancy prediction from whole slide cytopathology images (2019)
9. Folberg, R., Augsburger, J.J., Gamel, J.W., Shields, J.A., Lang, W.R.: Fine-needle aspirates of uveal melanomas and prognosis. *Am. J. Ophthalmol.* **100**(5), 654–657 (1985)
10. Garud, H., et al.: High-magnification multi-views based classification of breast fine needle aspiration cytology cell samples using fusion of decisions from deep convolutional networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 828–833, July 2017
11. Girard, N., Zhygallo, A., Tarabalka, Y.: ClusterNet: unsupervised generic feature learning for fast interactive satellite image segmentation. In: *Image and Signal Processing for Remote Sensing XXV*, vol. 11155, pp. 244–254. SPIE (2019)
12. Grossmann, P., Stringfield, O., El-Hachem, N., et al.: Defining the biological basis of radiomic phenotypes in lung cancer. In: *eLife* (2017)
13. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
14. Li, J., Li, W., Gertych, A., Knudsen, B.S., Speier, W., Arnold, C.W.: An attention-based multi-resolution model for prostate whole slide imageclassification and localization. *CoRR abs/1905.13208* (2019)
15. Lin, H., Chen, H., Dou, Q., Wang, L., Qin, J., Heng, P.: ScanNet: a fast and dense scanning framework for metastatic breast cancer detection from whole-slide image. In: *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 539–546 (2018)
16. Liu, J., Chen, F., Shi, H., Liao, H.: Single image super-resolution for MRI using a coarse-to-fine network. In: Ibrahim, F., Usman, J., Ahmad, M.Y., Hamzah, N., Teh, S.J. (eds.) *ICIBEL 2017*. IP, vol. 67, pp. 241–245. Springer, Singapore (2018). https://doi.org/10.1007/978-981-10-7554-4_42
17. Liu, T.A., Zhu, H., Chen, H., et al.: Gene expression profile prediction in uveal melanoma using deep learning: a pilot study for development of an alternative survival prediction tool. *Ophthalmol. Retina* **S2468-6530**, 30251–30257 (2020)
18. Paszke, A., Gross, S., Chintala, S., et al.: Automatic differentiation in PyTorch (2017)
19. Roullier, V., Lézoray, O., Ta, V.T., Elmoataz, A.: Multi-resolution graph-based analysis of histopathological whole slide images: application to mitotic cell extraction and visualization. *Computer. Med. Imaging Graph.* **35**(7), 603–615 (2011). *Whole Slide Image Process*

20. Saikia, A.R., Bora, K., Mahanta, L.B., Das, A.K.: Comparative assessment of CNN architectures for classification of breast FNAC images. *Tissue Cell* **57**, 8–14 (2019). eM in cell and tissues
21. Singh, A.D., Turell, M.E., Topham, A.K.: Uveal melanoma: Trends in incidence, treatment, and survival. *Ophthalmology* **118**(9), 1881–1885 (2011)
22. Wang, G., Li, W., Zuluaga, M.A., et al.: Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE Trans. Med. Imaging* **37**(7), 1562–1573 (2018)
23. Xu, N., Price, B., Cohen, S., Yang, J., Huang, T.S.: Deep interactive object selection. In: *IEEE Conference on Computer Vision and Pattern Recognition*, June 2016
24. Yang, B., Fu, X., Sidiropoulos, N., Hong, M.: Towards k-means-friendly spaces: simultaneous deep learning and clustering. In: *34th International Conference on Machine Learning, ICML 2017*, pp. 5888–5901 (2017)
25. Yang, B., Fu, X., Sidiropoulos, N., Hong, M.: Towards k-means-friendly spaces: simultaneous deep learning and clustering. In: *34th International Conference on Machine Learning, ICML 2017*, pp. 5888–5901 (2017)
26. Zhu, Z., Xia, Y., Shen, W., Fishman, E., Yuille, A.: A 3D coarse-to-fine framework for volumetric medical image segmentation. In: *2018 International Conference on 3D Vision (3DV)*, pp. 682–690, September 2018
27. Zhu, Z., Xia, Y., Xie, L., Fishman, E.K., Yuille, A.L.: Multi-scale coarse-to-fine segmentation for screening pancreatic ductal adenocarcinoma. *CoRR* abs/1807.02941 (2018)