

Learning from Similarity Proportion Loss for Classifying Skeletal Muscle Recovery Stages

Yu Yamaoka¹, Weng Ian Chan¹, Shigeto Seno¹,
Soichiro Fukada², Hideo Matsuda¹

¹ Graduate School of Information Science and Technology,
Osaka University, Osaka, Japan

{yu-yamaoka, chan.wengian, senoo}@ist.osaka-u.ac.jp

² Graduate School of Pharmaceutical Sciences, Osaka University, Osaka, Japan

Abstract. Evaluating the regeneration process of damaged muscle tissue is a fundamental analysis in muscle research to measure experimental effect sizes and uncover mechanisms behind muscle weakness due to aging and disease. The conventional approach to assessing muscle tissue regeneration involves whole-slide imaging and expert visual inspection of the recovery stages based on the morphological information of cells and fibers. There is a need to replace these tasks with automated methods incorporating machine learning techniques to ensure a quantitative and objective analysis. Given the limited availability of fully labeled data, a possible approach is Learning from Label Proportions (LLP), a weakly supervised learning method using class label proportions. However, current LLP methods have two limitations: (1) they cannot adapt the feature extractor for muscle tissues, and (2) they treat the classes representing recovery stages and cell morphological changes as nominal, resulting in the loss of ordinal information. To address these issues, we propose Ordinal Scale Learning from Similarity Proportion (OSLSP), which uses a similarity proportion loss derived from two bag combinations. OSLSP can update the feature extractor by using class proportion attention to the ordinal scale of the class. Our model with OSLSP outperforms large-scale pre-trained and fine-tuning models in classification tasks of skeletal muscle recovery stages.

Keywords: LLP · Skeletal muscle · WSI · Weakly supervised learning

1 Introduction

Evaluating the regeneration process of damaged muscle tissues is a fundamental analysis for measuring the effect size of biological experimental manipulations for discovering mechanisms of muscle weakness associated with aging and disease. One method to induce damage and recovery of muscle tissues is an injection of cardiotoxin (CTX), where CTX is injected into mice’s lower leg muscles

Yu Yamaoka and Weng Ian Chan contributed equally to this work.

to cause local necrosis of myofibers within the tissue. In recovery, damaged-myofiber-derived factors activate muscle satellite cells, which proliferate and become myoblasts. Many myoblasts fuse to form myotubes, and as they mature, they become embedded in the tissue and return to myofibers. The recovery speed of cells and fibers in muscle tissue exhibits locality [17], requiring the evaluation of the condition of cells in each specific region to assess tissue recovery [4]. This process is labor-intensive because muscle tissue images are high-resolution whole slide images (WSI) containing numerous cells and fibers. Therefore, effective and objective automated image analysis is necessitated.

Although supervised learning, as represented by deep learning, is a useful method, annotating recovery stages in every region of WSI is similarly labor-intensive and requires specialized expertise. In this context, the classification of multiple instances by learning label proportion (LLP) is one of the leading methods of weakly supervised learning in WSI analysis because it avoids GPU memory limitations caused by the large size of WSI and enhances the explainability of image analysis by predicting each region of the WSI. Examples of applying LLP to WSI include binary [19] or three-class [8] necrosis determination in tumors and screening for regions of interest [3]. LLP currently faces two challenges in evaluating the cell regeneration process. (1) As LLP computes the loss function after predicting all instances in a bag, it typically uses pre-trained models as feature extractors (backbones) and only updates the classification layer (head) [19, 8]. Large pre-trained models such as DINO [2] are generally trained on common objects, so they may not perform well in extracting features for medical WSI tasks. (2) Cells in the muscle recovery process gradually change their morphology over time, creating a natural biological order among the stages, which needs to be adequately captured. Specifically, the process involves the following stages: intact myofiber, ghost fibers (which are the basal membrane remnants post-CTX injection), myotubes (satellite cells that have undergone myogenesis), and recovered myofibers (cells that have fully regenerated). However, existing LLP methods treat the process on a nominal scale, leading to missing ordinal information. Therefore, we aim to address both issues by updating the feature extractor while considering the order of the classes in the training process under the LLP paradigm.

In this study, we propose an ordinal scale learning from similarity proportion: OSLSP that uses the date elapsed from CTX injection as weak supervision tied to cell class proportions, viewing morphological changes in cells during recovery as ordinal classes. A similarity proportion loss in OSLSP simultaneously addresses the issues of non-updated feature extractors in LLP and the lack of ordinal information among classes. Similarity proportion loss is computed between two bags, where the ground truth similarity proportion of classes is derived from the binomial theorem based on the class proportions of the two bags. During training, the similarity of feature vectors of instances from different bags is calculated, forming a similar distribution when comparing instances from two bags, and the loss function is computed by comparing this to the ground truth proportion. The

feature extractor trained with OSLSP contributed to improved accuracy in cell classification tasks in the skeletal muscle recovery dataset.

2 Related work

The weakly-supervised [14, 10, 15, 11] methods to address the challenges of WSI are shown in Table.1. Previous studies have designed loss functions using class proportions for bag-level learning and generated pseudo-labels for each instance from class proportions for instance-level learning[11]. While instance-level learning allows for the updating of feature extractors, it has been pointed out that pseudo-labels can lead to noise due to incorrect label generation [11]. Additionally, previous bag-level learning approaches have treated the output results as nominal scales, focusing on binary tumor classification [11, 14, 15], ternary classification [8], and the extraction of regions of interest in pathology images [3]. They have not targeted morphology changes of cells with ordinal scale classes in muscle tissues. Although not targeting WSI, Jerónimo’s [6] used already quantified patient data and implantation rates for embryo implantation prediction. LLP-VAT [18] proposes a consistency loss to ensure class prediction consistency even with slight noise added to individual instance images, in addition to the traditional LLP loss. SIM-LLP [7] added a pairwise similarity-based loss that penalizes different predictions of feature vectors in addition to the conventional LLP loss. Previous LLP methods have used pre-trained encoders for feature extraction, focusing on learning the classification head. This study aims to update the feature extraction backbone using similarity proportion loss.

Table 1. This table lists methods for addressing WSI classification tasks. *Class Scale* indicates the output scale by a classifier head. *Learning Target* denotes the layers of the model updated, where *Backbone* refers to the layer responsible for feature extraction, and *Head* refers to the layer that estimates the class from the extracted features. *Loss* indicates the loss function used for learning. *Prop.* stands for proportion.

Method	Class Scale	Learning Target	Loss
Pseudo [14]	Nominal	Head	Pseudo Class Label + Class Prop.
LLP [10]	Nominal	Head	Class Prop.
MIL [15]	Binary	Head	Class Prop.
IIB-MIL [11]	Binary	Backbone + Head	Pseudo Class Label + Class Prop.
OSLSP (Ours)	Ordinal	Backbone + Head	Similarity Prop. + Class Prop.

3 Method

3.1 OSLSP: Ordinal scale learning from similarity proportion

We propose a similarity proportion loss to simultaneously address the issues of updating the feature extractor and missing ordinal class information in LLP. As




Fig. 1. Overview of our OSLSP. **(a)** Similarity of class and morphological changes of cells over time. **(b)** A pipeline to obtain instance features and class inference results for each cell from WSI. **(c)** We computed similarity proportion loss using KL divergence between ground truth similarity distribution derived by combining two bags’ proportions and predicted one derived by cosine similarity of each instance in two bags.

shown in Fig. 1(a), we define classes $C = 1, \dots, k, \dots, K$ for different cell morphologies: Red: Intact Myofiber (Intact MF), Blue: Ghost fiber, Yellow: Myoblast, Orange: Myotube, Pink: Recovered Myofiber (Recovered MF). The dates after CTX injection are denoted as $d \in \{\text{“day0”}, \text{“day3”}, \text{“day5”}, \text{“day7”}, \text{“day14”}\}$, where Day 0 represents the state before CTX injection, meaning 100% cells are Intact MF. Ghost Fiber is the cell stage where myofibers transition immediately after CTX injection, with the highest proportion on Day 3. Over time, the proportion shifts towards myoblasts and myotubes as days pass. The necrotic cells lose all morphology of the ex-myofiber as shown in Fig. 1(a) “Myoblast”. Therefore, we see myoblasts as the birth of myofibers, while ghost fibers represent the end of them, making their similarity the most distant, as illustrated in Fig. 1(a) “sim”. At the stage where myoblasts have sufficiently proliferated, they become multinucleated myotubes, which means from Day 5 to Day 7, the proportion of myotubes increases. As myotubes grow, they fill the gaps between cells, resulting in Recovered MF. From this prior knowledge and observation of cell morphology, we perform rough manual class annotation on the training data, obtaining proportions $\mathbf{p}_d \in [0, 1]^K$ relative to each date label d .

After dividing the WSI into 256×256 [pixel] patch images, we use Cellpose [16] for cell segmentation and clip 64×64 [pixel] images with each segmented cell in the center for inference, as shown in Fig. 1(b). We obtain instance feature vectors $\mathbf{x} \in \mathbb{R}^D$ ($D \in \mathbb{N}$: the dimension of the image feature map) from 64×64 [pixel] images through a feature extractor (backbone). The classifier model produces inference output $\mathcal{F} : \mathbb{R}^D \rightarrow [0, 1]^K$ ($\sum_{k=1}^K \mathcal{F}(\mathbf{x})_k = 1$), where $\mathcal{F}(\mathbf{x})_k$ represents the classifier's confidence in class k .

Following the standard bag-making method [1] of LLP, we group N instances $\mathbf{x} \in \mathbb{R}^D$ of the same date into one bag, which has a corresponding class proportion $\mathbf{p}_d \in [0, 1]^K$ ($\|\mathbf{p}_d\|_1 = 1$, $\|\cdot\|_1$ means L1 norm), N is defined as the bag size. As shown in Fig. 1(c), we computed the similarity proportion loss by calculating Kullback–Leibler divergence between ground truth probability density function (PDF) P and predicted PDF $\hat{P}$. P is calculated by combining the ground truth class proportion of two bags, and $\hat{P}$ is computed as the cosine similarity between instance features in two bags. Given b as the number of bins of the discretized PDFs, the similarity proportion loss is defined as

$$\mathcal{L}_{SimProp} = D_{KL}(\hat{P} \| P) = \sum_{i=1}^b \hat{P}(i) \log \left(\frac{\hat{P}(i)}{P(i)} \right) \quad (1)$$

We compute the predicted PDF $\hat{P}$ by randomly picking two bags and compare the instance pairs $\mathbf{x}_n, \mathbf{y}_n$ ($n \in \{1, 2, \dots, N\}$) of the two bags using a scaled cosine similarity $\text{CosSim}(\mathbf{x}_n, \mathbf{y}_n) : \mathbb{R}^D \times \mathbb{R}^D \rightarrow [0, 1] \in \mathbb{R}$, which is defined as below.

$$\text{CosSim}(\mathbf{x}_n, \mathbf{y}_n) = \frac{1}{2} \left(\frac{\mathbf{x}_n \cdot \mathbf{y}_n}{\|\mathbf{x}_n\| \cdot \|\mathbf{y}_n\|} + 1 \right) \in [0, 1] \quad (2)$$

We obtain the cosine similarities $\text{CosSim}(\mathbf{x}_n, \mathbf{y}_n)$ for all instance pairs and plot them as a histogram. The parameters max , min , and b are used to set the maximum and minimum values of the data and the number of bins in the histogram. Since the cosine similarity CosSim values range from 0 to 1, we set $max = 1$ and $min = 0$. We compute the width of the histogram bins $\Delta = \frac{max-min}{b}$. The value of the i -th bin of the histogram $\text{hist}(i)$ is determined as

$$\text{hist}(i) = \sum_{n=1}^N \mathbf{I}(\text{CosSim}(\mathbf{x}_n, \mathbf{y}_n) \in [n \times \Delta, (n+1) \times \Delta]) \quad (3)$$

where $\mathbf{I}$ is an indicator function that counts values within the specified range. When we use the indicator function $\mathbf{I}$ in the design of the loss function, it results in a non-differentiable computation process, preventing the backpropagation of gradients from the loss to the model. We convert inter-atomic distances from discrete to continuous [13] to address this issue. We adopt Gaussian expansion with $\sigma = 0.1$ as described below instead of using Eq. 3.

$$\text{hist}(i) \cong \hat{P}(i) = \sum_{n=1}^N \frac{1}{\sigma\sqrt{2\pi}} \exp \left(-\frac{(\text{data}(n) - \mu_i)^2}{2\sigma^2} \right) \Delta \quad (4)$$

The histogram calculation is approximated in Gaussian expansion by integrating the Gaussian density functions for each data point $i \in \{1, 2, \dots, b\}$ over the bin width. Note that μ_i represents the mean similarity value of the i -th bin. This approach allows for a differentiable histogram calculation, enabling the gradient backpropagation necessary for training the model.

For any two classes $k, k' \in C = \{1, \dots, K\}$, we define a class similarity as described below for computing ground truth PDF P following the life cycle of myofibers, as shown in Fig. 1(a). When comparing the same class ($k = k'$), their similarity $\text{sim}(k, k')$ equals 1.

$$\text{sim}(k, k') = 1 - \frac{|k' - k|}{K - 1} \in [0, 1] \quad (5)$$

Using pair of the class proportions $\mathbf{p}_d, \mathbf{p}'_{d'} \in [0, 1]^K$ corresponding to date labels d and d' of the two bags, the ground truth PDF is calculated by

$$\begin{cases} P(\text{sim}(k, k')) = p_k \cdot p'_{k'} & \text{if } k = k'; \\ P(\text{sim}(k, k')) = p_k \cdot p'_{k'} + p_{k'} \cdot p'_k & \text{otherwise.} \end{cases} \quad (6)$$

3.2 LLP learning for classifier

To train a classifier head, we compute a proportion loss by calculating KL divergence between ground truth class proportions $\mathbf{p}_d \in [0, 1]^K$ and predicted one $\hat{\mathbf{p}}_d \in [0, 1]^K$ as below.

$$\mathcal{L}_{prop} = D_{\text{KL}}(\mathbf{p}_d \parallel \hat{\mathbf{p}}_d) = \sum_{k=1}^K p_d \log \left(\frac{p_d}{\hat{p}_d} \right) \quad (7)$$

$\mathcal{L}_{prop}$ is computed for each bag, which have $\mathbb{B} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ and class proportion $\mathbf{p}_d \in [0, 1]^K$ ($\|\mathbf{p}_d\|_1 = 1$) corresponding date labels d . Through the classifier $\mathcal{F} : \mathbb{R}^D \rightarrow \mathbb{R}_+^K$, we obtain the class confidence $\mathcal{F}(\mathbf{x}) \in [0, 1]^K$ and we obtain the predicted class proportions $\hat{\mathbf{p}}_d = [\hat{p}_1, \dots, \hat{p}_k, \dots, \hat{p}_K] \in [0, 1]^K$, $\|\hat{\mathbf{p}}_d\|_1 = 1$ by aggregating the class results of each instance at the bag level as below.

$$\hat{p}_k = \frac{1}{|\mathbb{B}|} \sum_{\mathbf{x} \in \mathbb{B}} \mathcal{F}(\mathbf{x})_k \quad (8)$$

4 Evaluation

Experiment settings. In our evaluation, we use 31 WSIs and obtain ground truth by professional annotation into 5 classes, one WSI per day. Thus, we have 26 WSIs for training and 5 WSIs for testing. Please refer to the supplementary materials for more details. We augment the data by rotation, flipping, and simulation of random optical conditions, including RandomBrightness (p=0.5), RandomContrast (p=0.5), and RandomGamma (p=0.5) on albumentations v1.3.1 after cutting to 256×256 [pixel]. To correctly reflect the proportion in the

dataset, we randomly clip 64×64 images from the 256×256 pixels images for training data. We use ViT-B/8 model³, a vision transformer model with a patch size of 8 as the backbone. We use the student checkpoint and enable the average pooling patch tokens. Other parameters are kept in default. In OSLSP, we only fine-tune the last block of the model. In DINO fine-tuning, we follow the official training pipeline and start the training from the pre-trained ViT-B/8 model. In training a classifier head by LLP using the class proportions of dates after preparing the feature extractor, we use a three-layer perceptron [12] with ReLU [9] activations. During inference, 64×64 images were clipped from the test data centered around cells based on segmentation obtained using a Cellpose [16] model, fine-tuned with training data on top of the Cyto model. We compare our OSLSP with a pre-trained DINO model [2] and a DINO model [2] fine-tuned with our training data.

Classification results. As shown in Table 2, we compare OSLSP with the baselines using manual expert annotations as ground truth. RMSE reflects the weight of incorrect predictions in class order “blue, red, pink, orange, yellow” as in Fig. 1(a). Fig. 2 shows the visual result of OSLSP and our baselines. The empty spaces of Manual in Fig. 2 indicate areas where even experts could not indeterminate cell class. When performing the quantitative evaluation of Table 2, the cells in this area are excluded from the calculations.

5 Discussions and Conclusion

Discussions. Applying the concept of contrastive learning to LLP enables the updating of feature extractors using class proportion as weak supervision. OSLSP reduces the domain gap of images between training and test data more effectively than the simple fine-tuned DINO model and better captures skeletal muscle cell features. In the supplementary, we include the UMAP plots for all feature extractors for better illustrations. As a result, we observe improvements in both RMSE and accuracy for OSLSP in Table 2, which indicates that OSLSP learns the order of the classes more correctly. In contrast, although the fine-tuned model exhibits a low RMSE, it converges to predicting mostly intact MF, recovered MF, and myotubes (as in Fig. 2), significantly harming the accuracy.

³ Self-Supervised Vision Transformers with DINO. <<https://github.com/facebookresearch/dino>>, last accessed on June 29, 2024.

Table 2. Classification results of LLP classifiers. Accuracy, Micro precision, recall, and F1-score. RMSE: Root mean square error based on ordinal cell class.

Method	Accuracy [%] $\uparrow$	Recall $\uparrow$	Precision $\uparrow$	F1-score $\uparrow$	RMSE $\downarrow$
DINO Pre-trained [2]	44.442	0.436	0.332	0.377	2.431
DINO Fine-tuned [2]	20.967	0.251	0.191	0.217	1.76
OSLSP (Ours)	46.005	0.492	0.375	0.425	2.152




Fig. 2. Classification results of WSIs for each day. Blue: ghost fiber, Red: intact myofiber, Pink: recovered myofiber, Orange: myotube, Yellow: myoblast. In the manual expert annotations, unannotated white areas indicate uncertain regions.

Unlike traditional LLP that treats classes as nominal scales, OSLSP considers the ordinal scale nature of class similarity, given that intact MF and ghost fibers are in similar states immediately before and after CTX injection. Our approach is similar to applying class weighting in supervised learning [5], which assigns different levels of importance to errors in the estimation during training. One of our limitations, as shown in Fig. 2, is that intact MF and ghost fiber are more frequently confused than the pre-trained model due to the high pre-set similarity of the two classes. Care must be taken with the weighting between these classes,

as intact MF and necrotic ghost fibers, while morphologically similar, are totally different regarding muscle strength. Therefore, misclassifying the two classes might have a substantial negative impact when measuring muscle strength.

Conclusion. In this study, we propose that OSLSP update a feature extractor by a similarity proportion loss, allowing the injection of prior knowledge into class order. OSLSP enables ordinal scale learning tailored to the objective and achieves better analysis of muscle recovery.

6 Acknowledgements

JSPS KAKENHI Grant Numbers JP22H05085 and JP22K12246 partially supported this work.

References

1. Brahmbhatt, A., Pokala, M., Saket, R., Raghuveer, A.: Llp-bench: A large scale tabular benchmark for learning from label proportions. arXiv preprint arXiv:2310.10096 (2023). <https://doi.org/10.48550/arXiv.2310.10096>
2. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
3. Chan, L., Hosseini, M.S., Rowsell, C., Plataniotis, K.N., Damaskinos, S.: Histoseg-net: Semantic segmentation of histological tissue type in whole slide images. In: International Conference on Computer Vision (ICCV). pp. 10661–10670 (2019). <https://doi.org/10.1109/ICCV.2019.01076>
4. Desgeorges, T., Liot, S., Lyon, S., Bouvière, J., Kemmel, A., Trignol, A., Rousseau, D., Chapuis, B., Gondin, J., Mounier, R., Chazaud, B., Juban, G.: Open-CSAM, a new tool for semi-automated analysis of myofiber cross-sectional area in regenerating adult skeletal muscle. *Skeletal Muscle* **9** (2019). <https://doi.org/10.1186/s13395-018-0186-6>
5. Gao, B.B., Xing, C., Xie, C.W., Wu, J., Geng, X.: Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing* **26**(6), 2825–2838 (2017). <https://doi.org/10.1109/TIP.2017.2689998>
6. Hernández-González, J., Inza, I., Crisol-Ortiz, L., Gumenbe, M., Iñarra, M., Lozano, J.: Fitting the data from embryo implantation prediction: Learning from label proportions. *Statistical Methods in Medical Research* **27** (05 2016). <https://doi.org/10.1177/0962280216651098>
7. Kotzias, D., Denil, M., de Freitas, N., Smyth, P.: From group to individual labels using deep features. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. p. 597–606 (2015). <https://doi.org/10.1145/2783258.2783380>
8. Matsuo, S., Suehiro, D., Uchida, S., Ito, H., Terada, K., Yoshizawa, A., Bise, R.: Learning from partial label proportions for whole slide image segmentation. In: Medical Image Computing and Computer Assisted Intervention (MICCAI) (2024). <https://doi.org/10.48550/arXiv.2405.09041>
9. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: Proceedings of International Conference on Machine Learning. p. 807–814 (2010)

10. Quadrianto, N., Smola, A.J., Caetano, T.S., Le, Q.V.: Estimating labels from label proportions. In: Proceedings of the 25th International Conference on Machine Learning. p. 776–783. ICML '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1390156.1390254>
11. Ren, Q., Zhao, Y., He, B., Wu, B., Mai, S., Xu, F., Huang, Y., He, Y., Huang, J., Yao, J.: lib-mil: Integrated instance-level and bag-level multiple instances learning with label disambiguation for pathological image analysis. In: Medical Image Computing and Computer Assisted Intervention (MICCAI). pp. 560–569 (2023). https://doi.org/10.1007/978-3-031-43987-2_54
12. Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review* **65**(6), 386 (1958). <https://doi.org/10.1037/h0042519>
13. Schütt, K.T., Arbabzadah, F., Chmiela, S., Müller, K.R., Tkatchenko, A.: Quantum-chemical insights from deep tensor neural networks. *Nature communications* **8**(1), 13890 (2017). <https://doi.org/10.1038/ncomms13890>
14. Shao, Z., Wang, Y., Chen, Y., Bian, H., Liu, S., Wang, H., Zhang, Y.: Lnpil-mil: Learning from noisy pseudo labels for promoting multiple instance learning in whole slide image. In: International Conference on Computer Vision (ICCV). pp. 21438–21438 (2023). <https://doi.org/10.1109/ICCV51070.2023.01965>
15. Shao, Z., Bian, H., Chen, Y., Wang, Y., Zhang, J., Ji, X., yongbing, z.: Transmil: Transformer based correlated multiple instance learning for whole slide image classification. In: Advances in Neural Information Processing Systems (2021). <https://doi.org/10.48550/arXiv.2106.00908>
16. Stringer, C., Wang, T., Michaelos, M., Pachitariu, M.: Cellpose: a generalist algorithm for cellular segmentation. *Nature Methods* **18**(1), 100–106 (2021). <https://doi.org/10.1038/s41592-020-01018-x>
17. Tanaka, Y., Kita, S., Nishizawa, H., Fukuda, S., Fujishima, Y., Obata, Y., Nagao, H., Masuda, S., Nakamura, Y., Shimizu, Y., Ryohei, M., Tomoaki, N., Tohru, F., Barbara, R., So-Ichiro, F., Norikazu, M., Iichiro, S.: Adiponectin promotes muscle regeneration through binding to t-cadherin. *Scientific Reports* **9**(1), 16 (2019). <https://doi.org/10.1038/s41598-020-66545-1>
18. Tsai, K.H., Lin, H.T.: Learning from label proportions with consistency regularization. In: Asian Conference on Machine Learning. pp. 513–528. PMLR (2020). <https://doi.org/10.48550/arXiv.1910.13188>
19. Ye, Q., Zhang, Q., Tian, Y., Zhou, T., Ge, H., Wu, J., Lu, N., Bai, X., Liang, T., Li, J.: Method of tumor pathological micronecrosis quantification via deep learning from label fuzzy proportions. *IEEE Journal of Biomedical and Health Informatics* **25**(9), 3288–3299 (2021). <https://doi.org/10.1109/jbhi.2021.3071276>