

Hierarchical Semantic Segmentation using Psychometric Learning

Lu Yin¹

Vlado Menkovski¹

Shiwei Liu¹

Mykola Pechenizkiy^{1,2}

L.YIN@TUE.NL

V.MENKOVSKI@TUE.NL

S.LIU3@TUE.NL

M.PECHENIZKIY@TUE.NL

¹ *Eindhoven University of Technology, Eindhoven 5600 MB, Netherlands*

² *University of Jyväskylä, P.O. Box 35, FI-40014, Finland*

Abstract

Assigning meaning to parts of image data is the goal of semantic image segmentation. Machine learning methods, specifically supervised learning is commonly used in a variety of tasks formulated as semantic segmentation. One of the major challenges in the supervised learning approaches is expressing and collecting the rich knowledge that experts have with respect to the meaning present in the image data. Towards this, typically a fixed set of labels is specified and experts are tasked with annotating the pixels, patches or segments in the images with the given labels. In general, however, the set of classes does not fully capture the rich semantic information present in the images. For example, in medical imaging such as histology images, the different parts of cells could be grouped and sub-grouped based on the expertise of the pathologist.

To achieve such a precise semantic representation of the concepts in the image, we need access to the full depth of knowledge of the annotator. In this work, we develop a novel approach to collect segmentation annotations from experts based on psychometric testing. Our method consists of the psychometric testing procedure, active query selection, query enhancement, and a deep metric learning model to achieve a patch-level image embedding that allows for semantic segmentation of images. We show the merits of our method with evaluation on the synthetically generated image, aerial image and histology image.

Keywords: Psychometric test, Hierarchical Semantic Segmentation, Deep metric learning, Active learning.

1. Introduction

One of the major challenges in the task of image semantic segmentation is collecting high-quality supervision from the domain expert. Typically this task is reduced to assign a label to an image region from a fixed set of labels (Vernaza and Chandraker, 2017; Wang et al., 2020; Khoreva et al., 2017). While in some cases, such reduction may have sufficient semantic fidelity in general, this is still a major limitation. One can imagine that in the majority of applications, the concepts present in the images have a much more complex relationship between each other and possibly form some kind of hierarchies. For example, when dealing with histology image segmentation, rather than forcing the expert to project their domain knowledge to flat labels (e.g., in the case of biopsy segmentation, one set of classes may include: normal, benign, in-situ and invasive (Aresta et al., 2019)), we would

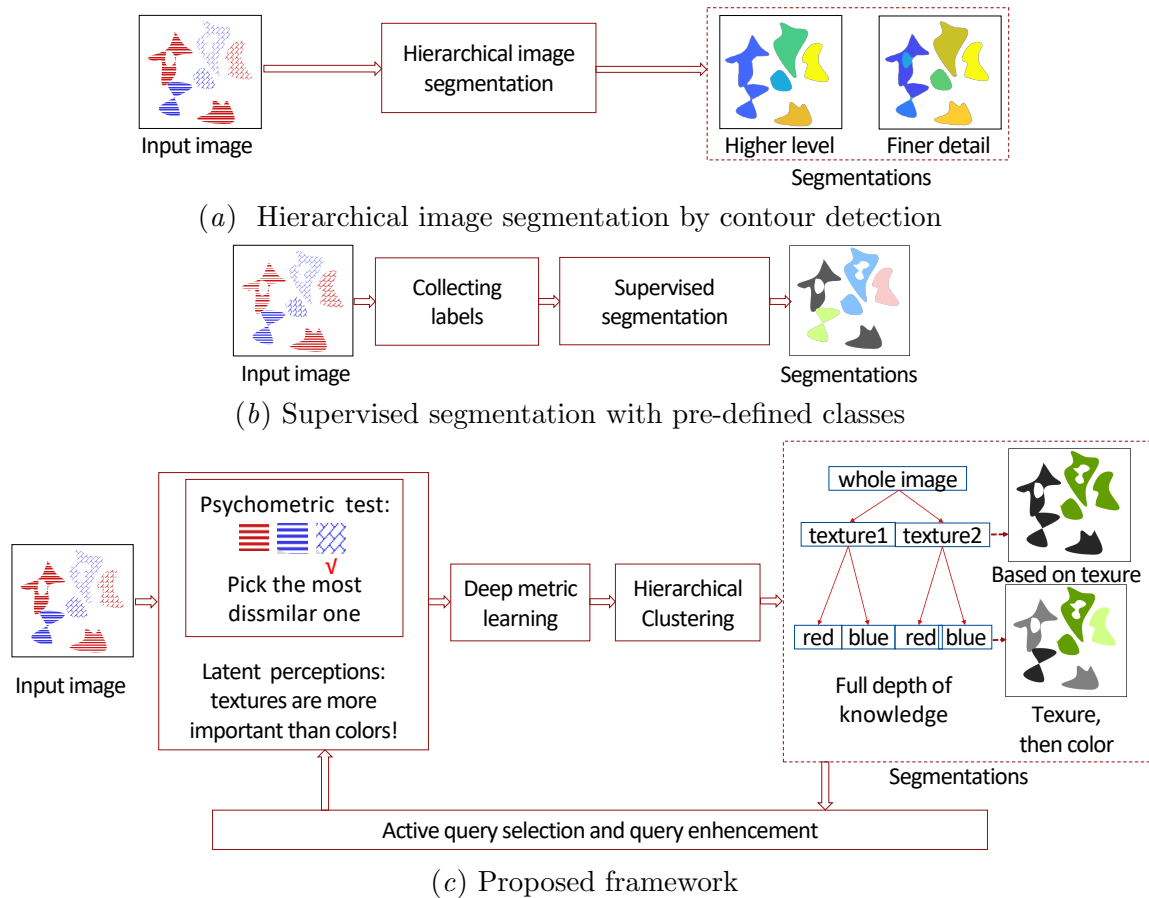


Figure 1: A comparison of three different segmentation methods using an example image containing six regions with two different textures and two different colors. (a): *Hierarchical image segmentation by contour detection*. Segmentation is implemented only based on information present in the image. Domain knowledge can not be utilized to form complex assignments. (b): *Supervised segmentation*. A fixed set of labels is assigned to each pixel. Coarse-level domain knowledge is expressed, but high fidelity hierarchical relationships present in the image are discarded. (c): *Our proposed framework*. Psychometric tests are applied to collect full-depth knowledge from the domain experts. In the test, three patches are sampled from the image for comparison. We assume the domain experts tend to choose the third option (blue cross slash) as the most dissimilar one since they perceive textures as more important than colors. After collecting a number of test responses, the image could be hierarchically segmented based on the texture on a higher level and then divided by colors, which reflects the captured experts’ perception structure. Besides, concept relationships are mapped to color similarities for visually inspecting. For instance, in the second level segmentation, two different but similar colors, dark green and light green, represent two also similar but not the same concepts, red cross slash and blue cross slash. Note that the square patches in the test are just for illustration and are replaced by super-pixels in the following experiments.

let them express their full-depth knowledge such that the experts can express the subtle differences in the contents of any two patches. Some scholars tried to use hierarchical segmentation method (Arbelaez et al., 2010) to capture the pattern relationships in an image. As no semantic supervision is provided from the domain expert, the built hierarchical structure is still not able to match the full-depth expert knowledge.

To address these challenges, we developed a psychometric-test-based approach to elicit the expert’s high fidelity knowledge, and map it to semantic hierarchical segmentation. Firstly, we create super-pixel patches by the SLIC (Achanta et al., 2012) approach and use psychometric test to measure the patches’ perceived similarity to each other. Next, we apply a deep metric learning algorithm to project the perceived relative distances to an embedded space where we can structure the patches in a hierarchical fashion, thereby proceed to image hierarchical segmentation. The built hierarchy reflects the latent perception structure, which we aim to capture from the annotator. To visualize the structure, we assign the segmented regions different color overlays from a given palette. The distance of the colors in the palette corresponds to the distances between patches in the embedded space. This allows us to visually inspect not only the captured relationships between the concepts of the image in the embedded space but also concretely the image.

Our proposed approach is illustrated in a toy example and compared with two baselines in Fig 1. It is shown that our method could capture the full-depth knowledge from experts and form a hierarchical semantic segmentation accordingly.

The our paper is organized as follows: in section 2 we review the most related researches; section 3 explains the details of our approach; in section 4 we evaluate our experiments in three different scenarios, carry an ablation study, and compare our approach with two baselines quantitatively; at last, we conclude our work and contributions in section 5.

2. The Related Work

There are three research lines closely related to our work: psychometric test, semantic image segmentation and deep metric learning.

Psychometric Test Psychometric test studies the perceptual processes of psychical stimuli (Gescheider, 2013). It has a long researched history and wide applications in different areas (Son et al., 2006; Feng et al., 2014). Some scholars use discriminative-based psychometric tests to measure the slightest difference between two stimuli that a subject can detect. It could provide more accrue results with fewer noises than directly assigning quantified values (Jogan and Stocker, 2014). Especially, the two-alternative-force choice (2AFC) model was developed in Fechner (1860), in which subjects are asked to compare the differences between two perceived stimuli and forced to make a correct choice. For example, two cups of water are given, the participant is asked to compare which one is sweeter. This experiment has been adopted to measure the subjects’ perception of more complex multimedia content such as images or videos (Son et al., 2006; Feng et al., 2014). As there are only two options in the 2AFC test, limited perception is captured by a test and some of the questions might be ambiguous to subjects. Some scholars extend the compared objects to m , forming the M-AFC family DeCarlo (2012). In this work, we use a variation of 2AFC, the three-alternative-force choice(3AFC) test, in which three samples are compared to capture the annotator’s perception of images.

Semantic image segmentation Based on the type of annotation, image segmentation could be divided into the supervised and semi-supervised way. The former one required pixel-level class labels, which is laborious but achieving better performance. FCN is a milestone for supervised segmentation and inspired many subsequent researches (Long et al., 2015). It provided an end-to-end trainable model and available for the arbitrary size of inputs. Parsenet (Liu et al., 2015) extended FCN by considering the global context information. Atrous convolution was integrated into the Deeplab segmentation family (Chen et al., 2018) to enlarge the receptive field without increasing the computational cost. Attention-mechanisms-based segmentation was explored to assign different weights to objects in Fu et al. (2019). Weakly supervised semantic segmentation (WSSS), on the other hand, uses weak labels instead of pixel-level annotations to guide model training, try to achieve a balance between performance and annotation costs. Different forms of supervision have been explored: image-level classification labels (Wang et al., 2020), bounding boxes (Khoreva et al., 2017) or scribbles (Vernaza and Chandraker, 2017).

However, these researches still suffer from the limitation of reducing the expert’s high-quality supervision to a fixed set of labels. In this paper, we apply the psychometric learning based way to collect full-depth knowledge from experts.

Deep metric learning and informative pair selection Metric learning aims to learn a measure of distance between different data points and has experienced great progress due to the development of deep learning (Kaya and Bilge, 2019). Many approaches are developed by the comparison-based model. Contrastive loss was applied to force a margin distance between positive and negative pairs by making pair-wise comparison (Hadsell et al., 2006). triplet-net (Schroff et al., 2015), quadruplet-net (Chen et al., 2017) and N-pair-net (Sohn, 2016) extend the compared sample numbers in a loss-function to three, four and N . Though yielding promising results, it is impractical to exhaust all the possible pair combinations for model training. Different approaches were proposed to select nontrivial samples. Semi-hard was applied in Schroff et al. (2015) to select triplets that conflict with the loss function. Suh et al. (2019) use class-based strategy to mining sample in a coarse-to-fine fashion. The most related work is from Yin et al. (2020). The scholars developed a dual-triplet loss function in which only negative image needs to be specified, and they applied a Bayesian-based scheme to actively select triplet queries when an actual human participant is involved. We followed that work’s spirit and discussed the impacts of different margin settings. Besides, we extend the query selection method to the segmentation scenario and propose a query enhancement strategy to improve the query efficiency further.

3. Approach

As shown in Fig 1(c)subfigure, our psychometric learning based semantic segmentation framework consists of four parts, psychometric test, deep metric learning, hierarchical clustering, and active query selection with query enhancement.

First, we use psychometric tests to elicit the annotator’s perception. Detailly, we use discriminative-based psychometric testing to perceive relative differences among super-pixel patches created by SLIC (Achanta et al., 2012). Our rich semantic knowledge could be mapped to a model by performing simple comparison tasks without assigning specific labels to image patches.

Then, a deep metric learning method is applied to project the perceived relative distances to embedding distances. That is, if an annotator believes two image patches are more similar than others, their relative distance should also be closer in the embedding space. Similar to [Yin et al. \(2020\)](#). We use the dual-triplet loss to align with our psychometric tests.

Next, in the embedding space, we cluster and sub-cluster super-pixel patches, organize them in a hierarchical fashion to segment images at different levels. The built hierarchy reflects the captured full-depth knowledge from the annotator. Unlike the hierarchical image segmentation research in [Arbelaez et al. \(2010\)](#), our segmentation is created based on the annotator’s responses instead of pure image information. Therefore, we are able to create various hierarchical segmentation based on different annotators’ senses.

To visualize the segmentation structure, we assign the segmented regions different colors from a given palette. Since these colors are projected from embedding space that represents the concepts’ relationships, the concept similarities could be inspected by colors intuitively.

While the effectiveness of psychometric testing, there could be too many potential queries due to the combination. It is crucial to develop a query selection scheme to select a fraction of proper queries so that the psychometric testing could be carried in an efficient way. Our work uses a Bayesian-based method to select queries with high uncertainty (informative) and high utility (un-ambiguous). Besides, a query augmentation method is applied to simulate more responses by the answered ones.

3.1. Psychometric Test

Three-alternative-force choice (3AFC) ([DeCarlo, 2012](#)) tests are applied to elicit the annotator’s perception of an image. To be specific, first, we use SLIC to create super-pixel patches from an image by color similarity and spatial proximity. It is commonly used as a pre-processing step in segmentation algorithms ([Achanta et al., 2012](#)). Then we sample three from the super-pixels in every test, and ask the annotators to choose the most dissimilar one. By making this simple decision, perceptions of the patches’ relative similarities are elicited. Note that, in the 3AFC test, besides super-pixels themselves, we present their surrounding areas as well to take advantage of more global image information (see [Fig. 2](#)).

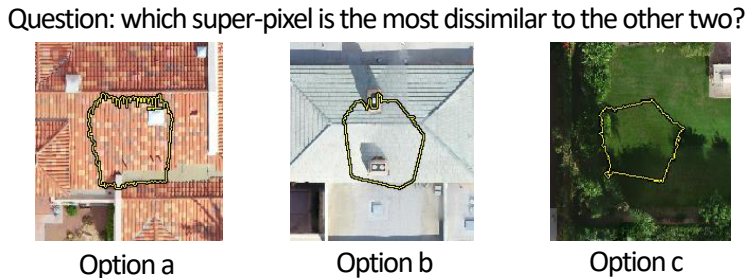


Figure 2: An example of three-alternative-forced-choice in the aerial image. The subject is asked to compare the super-pixel patches (patches circled by the yellow line) and choose the most dissimilar one against the other two. Surrounding areas are also presented for reference.

3.2. Deep Metric Learning

After psychometric testing, we use deep metric learning to map the captured patch similarities to embedding distances. There are varieties of deep metric learning models available, but a natural choice would be triplet-net (Schroff et al., 2015), as we also have three samples in the 3AFC test. The typical loss function in triplet-net is shown as,

$$L = \sum_{i=1}^N [\|f(x_a^i) - f(x_p^i)\| - \|f(x_a^i) - f(x_n^i)\| + m]_+ \quad (1)$$

where x_a^i, x_p^i are anchor and positive images from the same class. x_n^i is the negative image from a different class. $f(\cdot) \in \mathbb{R}^d$ is the learned embedding representation from a neural network. It embeds an image to a d-dimension euclidean space. $[\cdot]_+$ denotes a hinge function $\max[0, \cdot]$. This loss function enforce a distances margin m between images pairs with same labels to other image pairs.

Note that, in equation 1, anchor, negative and positive three images all have to be indicated. However, in a 3AFC test, only the most dissimilar one needs to be specified. To deal with this conflict, Yin et al. (2020) proposed a dual-triplet loss function which define the chosen image in 3AFC as negative x_n and rest two images all as positive images x_{p1}^i, x_{p2}^i . The dual-triplet loss introduces a dynamic margin to encourage the image pairs that the annotator perceives similarly to be closer than other pairs in embedding space. In our work, we use a constant margin m instead of a dynamic one and rewrite the dual-triplet loss as $L_{dual} =$

$$\sum_{i=1}^N \left\{ [\|f(x_{p1}^i) - f(x_{p2}^i)\| - \|f(x_{p1}^i) - f(x_n^i)\| + m]_+ + [\|f(x_{p1}^i) - f(x_{p2}^i)\| - \|f(x_{p2}^i) - f(x_n^i)\| + m]_+ \right\} \quad (2)$$

The reason for using a constant margin is that the elicited knowledge structure in segmentation is usually not that deep and complex, and a constant margin would be enough to build a discriminative embedding space without extra computation costs. Besides, the extra variables from by dynamic margin could introduce unstable factors.

3.3. Hierarchical Clustering

Once the semantic embedding is obtained, we could hierarchically cluster super-pixel patches in the embedding space to create hierarchical segmentation.

In general, there are two main categories of the hierarchical clustering algorithms, a bottom-up approach called ‘‘agglomerative’’ and a top-down approach called ‘‘divisive’’ (Rokach and Maimon, 2005). The agglomerative approach starts with each sample belonging to one single cluster and merges samples together as the hierarchy moves up. The divisive approach starts with all data as one cluster and splits recursively when the hierarchy goes down. Commonly this approach uses heuristics to choose splits, such as K-means.

In the agglomerative method, we need to carefully indicate the cutting value when we need a node containing more than two branches in the hierarchy. The cutting value is not intuitive compared with K in K-means because K could represent the number of branches directly. Therefore, in our work, we applied the divisive approach with K-means to create

hierarchical segmentation. Silhouette score (Kodinariya and Makwana, 2013) is applied to help us decide the cluster number K and when to stop divide.

We also map the concept relationships to color similarities for visually inspecting. In detail, we use MDS (Ghodsi, 2006) reducing the image’s embedding dimension to three and scale the embedding value to $0 \sim 255$ to represent RGB values. The concepts’ relative distances are then mapped to a color space. For instance, in a histopathological image, the color of in-situ areas should more like invasive areas than others if these two areas’ concepts are also more similar in the annotator’s perception.

3.4. Active Query Selection and Query Enhancement

In psychometric tests, there are $\binom{B}{3}$ potential questions given B patches by combination. Since it is impractical to answer all the queries and not every query contributes equal gradients for training, we developed a scheme to actively select proper queries with uncertainty and utility. Furthermore, we enhance the answered queries by simulating more responses.

Our selection procedure follows an iterative process. In each iteration, there are three steps. Our model and the elicited hierarchical knowledge will be updated by the selected queries in each iteration.

First, we generate potential questions by random sampling at each level of the hierarchy created from the last iteration. That allows the model to capture the global distribution as well as the local information in finer detail.

Secondly, inside the potential questions, we future select queries with high uncertainty. That is, queries that annotators are not confident with. Those queries are usually more informative and produce more training gradients (Settles, 2009). We use Bayes’ theorem to calculate the distribution of answering possibility and measure the uncertainty by the answers’ expected values.

To be specific, for a 3AFC query q_i , we denote the possibilities of picking three options a, b, c as $\theta = (\theta_1, \theta_2, \theta_3)$, and assume they follows dirichlet’s distribution $Dir(\alpha)$. The distribution density can be written as,

$$p(\theta|\alpha) = \frac{1}{Beta(\alpha)} \prod_{i=1}^3 \theta_i^{\alpha_i-1} (\theta_i \geq 0; \sum_{i=1}^3 \theta_i = 1), Beta(\alpha) = \frac{\prod_{i=1}^3 \Gamma(\alpha_i)}{\Gamma(\sum_{i=1}^3 \alpha_i)} \quad (3)$$

In the beginning, α is set to $(1, 1, 1)$, which means the possibility of choosing any option is the same. Along with the annotation process, we update the posterior distribution by the answered queries D . In the neighbor areas of q_i ’ each sampler, we choose the top k closest super-pixel patches by their embedding distances to the sample and consider the queries sampled from these patches as q_i ’s similar queries (see Fig 3 when $k = 1$). We seek q_i ’s similar queries in D , count the number of picking the similar queries’ three options as (m_1, m_2, m_3) , and update $Dir(\alpha)$ by $p(\theta|\alpha, D) = p(\theta|\alpha_1 + m_1, \alpha_2 + m_2, \alpha_3 + m_3)$. Then we calculate the expected value of choosing a, b, c by the updated posterior distribution. If any option’s expected picking possibility is higher than a threshold, we could conclude that it is a confident query for an annotator and reject this query.

The third step is to reject queries with less utility, which is assessed by the variance of a, b, c three options’ picking distribution. A query with high variance means it is still ambiguous to the annotator after he/she answered many similar queries and will be rejected.

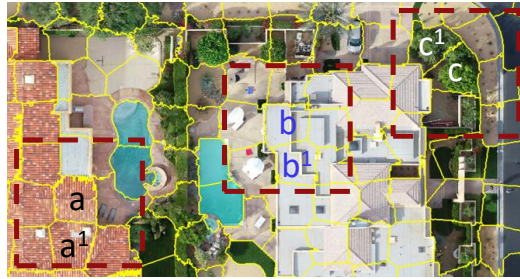


Figure 3: A query and its most similar query. a, b, c are three super-pixels from a query. In each sample’s neighbour area (bounded by red dotted line), we pick the closest super-pixels a_1, b_1, c_1 by their embedding distances to a, b, c . Query $[a_1, b_1, c_1]$ then is the similar query to $[a, b, c]$.

From Fig 3, we can notice that annotators tend to make the same decisions when facing similar queries. For example, in the query $[a, b, c]$, if the annotator consider the sample c as most dissimilar because b and a are all building patches, a similar choice of c_1 also likely to be made in query $[a_1, b_1, c_1]$. To further lighten the annotator’s burden, we simulate more answers by these similar queries.

4. Experiments

We conduct the experiments in three scenarios with real or virtual participants to show the value of our method.

We first evaluate our model in a synthetic image. As texture information is the main contributing factor in many scenarios, such as medical or aerial segmentation, a synthetic image with dominant texture information should be a good beginning to test our model. We simulate virtual annotators who always precisely respond based on given knowledge hierarchies to ensure no other variability is introduced, and compare the elicited knowledge structure with the given ones in an objective manner. To confirm that our method is able to create different hierarchical segmentation based on various perceptions, the simulated participants hold partial agreements about the latent perception.

Next, we test our method’s potentials in real-world applications. Two experts with different perceptions about a histopathological image are simulated for psychometric testing.

Then, we assess our model’s ability when dealing with real human participants with extra involved variations. The experiment is conducted on an aerial image.

We emphasize that the main contribution of our work is to elicit full-depth knowledge and concept relationships from experts rather than perfect pixel-wise prediction. Therefore, the commonly used mIoU score in image segmentation will not be regarded as a metric in this work. Instead, we use dendrogram purity (Heller and Ghahramani, 2005) to measure the matching degree between our extract hierarchical segmentation and the ground-truth perception, and use each cluster’s purity to quantify how accurate each segmented sub-area is. All the experiments use VGG19 as backbones. The margin value m is set to 0.2.

4.1. Virtual Participant on Synthetic Image

To remove other variability and test our method in a controllable way, we created a synthetic image that contains light green, normal green, dark green three colors, and lines, scatter dots, scatter triangles three textures. The nine different categories are shown in Fig 4. The whole image’s resolution is $1800 \times 3600 \times 3$ pixels.

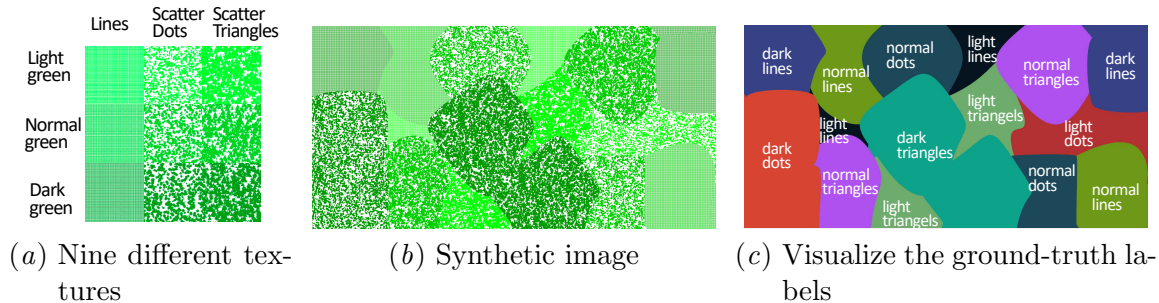


Figure 4: Synthetic image and the ground truth. We abbreviate dark green, normal green, light green to dark, normal, green and scatter dots, scatter triangles to dots, triangles respectively in (c) and the following figures in this paper.

Since human perception is not obvious before been elicited, it would be ambiguous to assess how well the extracted hierarchical segmentation matches the latent ground truth. Therefore, we create two virtual participants with given hierarchies to simulate the latent knowledge and generate responses accordingly.

To be specific, we create 1112 super-pixel patches by SLIC for psychometric tests and train the model with 10 iterations. Each iteration contains 800 simulated responses and another two times enhanced triplets. We illustrate the first participant’s segmentation result in Fig 5. The result shows that our model is capable of eliciting a clear knowledge structure that well matches the latent perception (simulated by given hierarchy), with all sub-cluster purity above 95%. In the first-layer segmentation, results are based on the areas’ color (dark green, normal green, light green). Then the segmentation splits into detailed nine categories in the last layer. Note that the color similarities match the concept relationships well in the segmentation visualization. For instance, all the normal green areas in the original image are marked as purple in Fig 5(d)subfigure as these areas sharing the same normal green color, but the purple has different shades since these areas have different textures (scatter triangles, scatter dots, lines). We can observe similar results on the second participant (see Fig 6), but with a different knowledge structure. The image is segmented based on textures on a higher layer then divided by colors. That confirms our model’s ability to create various hierarchical segmentation results based on different latent perceptions.

4.2. Virtual Participants on Histopathological Image

Then we evaluate our model on the real-world histopathological image from BACH Challenge 2018 (Aresta et al., 2019). We crop the image to $4211 \times 10020 \times 3$ pixels, create 2574 super-pixels, and train the model with 10 iterations. 1000 responses are simulated in each iteration except 1500 responses in the first one. Answered triplets are doubled by the

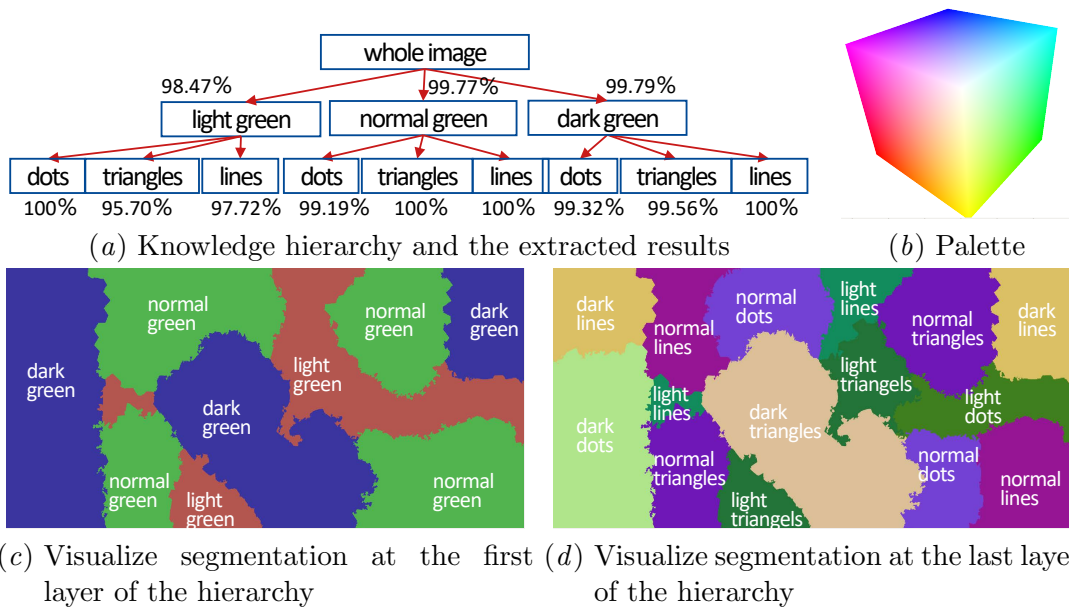


Figure 5: (a) is a given hierarchy to simulate the latent knowledge structure from virtual participant 1. The value by each node indicates the cluster’s purity extracted by our method. (b) is a palette for color similarity reference and could be used in following figures in the paper. (c) and (d) are segmentation results in the hierarchy’s first and last layer, and if their colors are similar based on the palette, the concepts that colors represent should also be considered close.

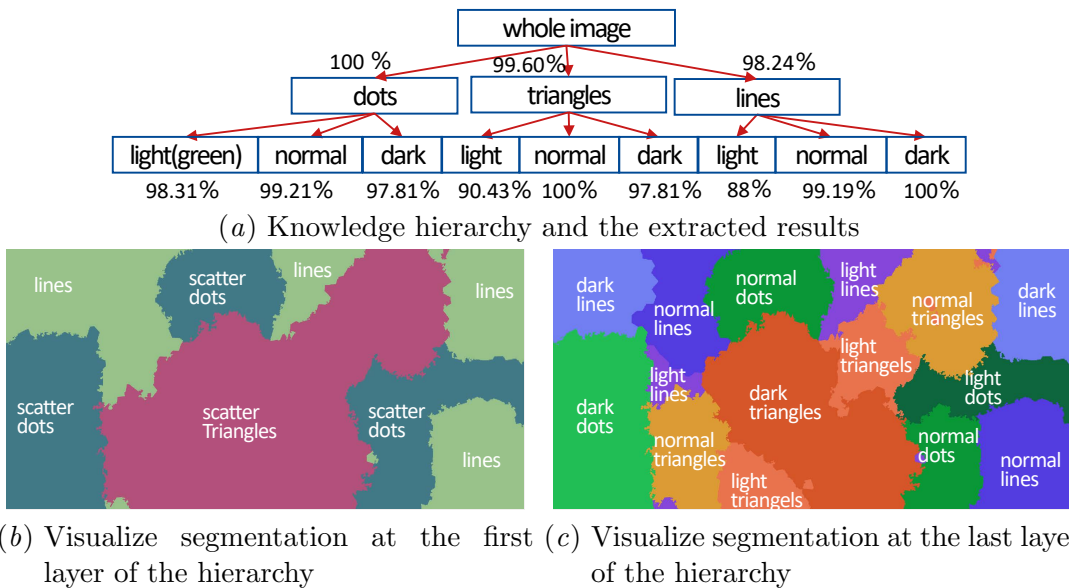


Figure 6: Knowledge structure from virtual participant 2 and hierarchical segmentation results on synthetic image.

proposed enhancement algorithm before used for training. Two participants with different latent hierarchical knowledge are simulated.

We present the elicited hierarchy and segmentation results of participant 1 in Fig 8. It can be seen that a clear hierarchical structure is extracted with each note’s purity higher than 94%. The first layer of segmentation is built according to if a cell is healthy. Then it goes down to four specific categories. Besides, the colors of benign, in-situ, invasive are similar to each other but have a big difference from the normal area in visualization Fig 8(c)subfigure, which matches the conceptual similarity in the elicited knowledge hierarchy. The results of participant 2 are shown in Fig 9. As the different knowledge tree from participant 1, the higher-level segmentation is formed based on if the area is cancer.

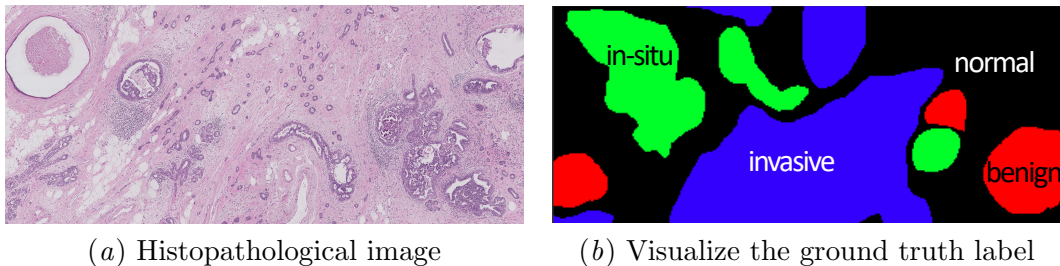


Figure 7: Histopathological image and the ground truth label.

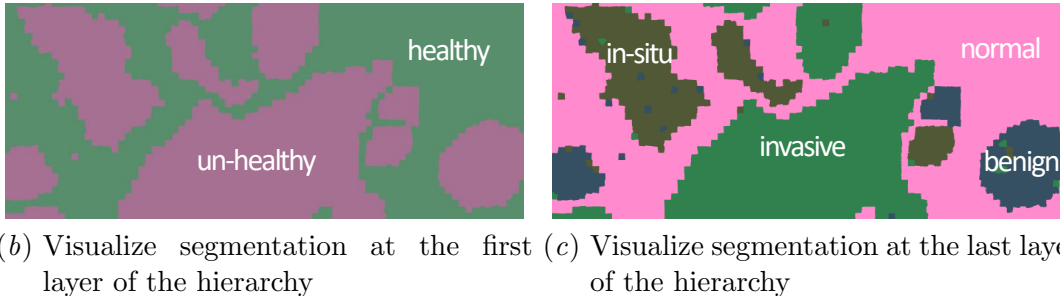
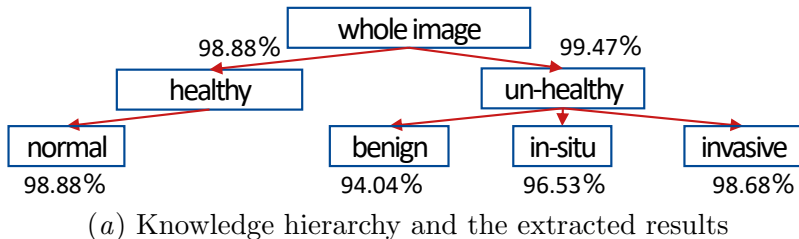


Figure 8: Segmentation results of virtual participant 1 on histopathological image.

4.3. Real Participant on Aerial Image

At last, we involved a real human in our psychometric tests. The goal is to test our model’s ability to deal with extra variability introduced by a human operator. 821 super-pixel patches are created from a $2128 \times 1619 \times 3$ aerial image¹ for psychometric tests. The

1. From repository <https://github.com/dronedeploy/dd-ml-segmentation-benchmark>

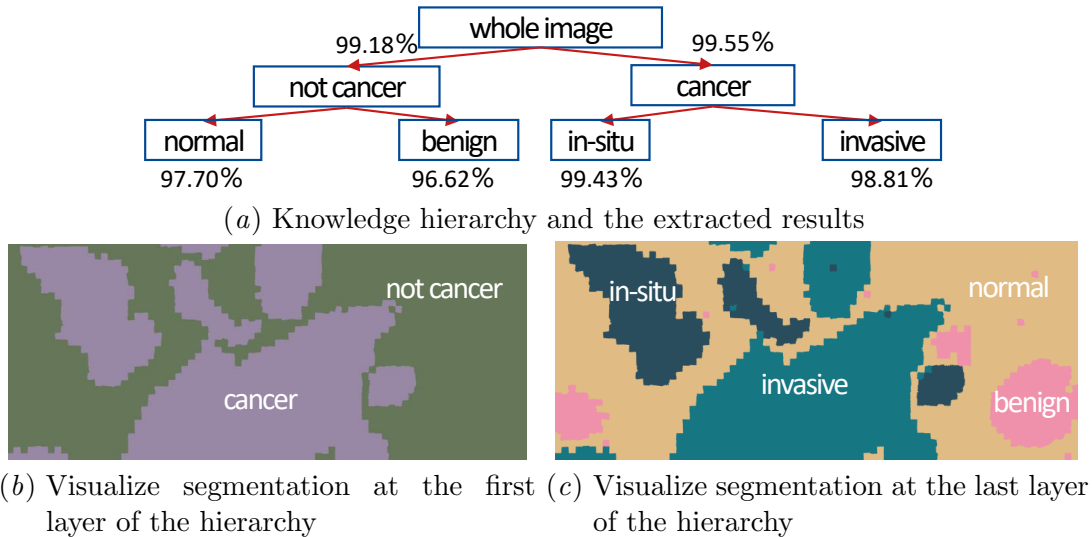


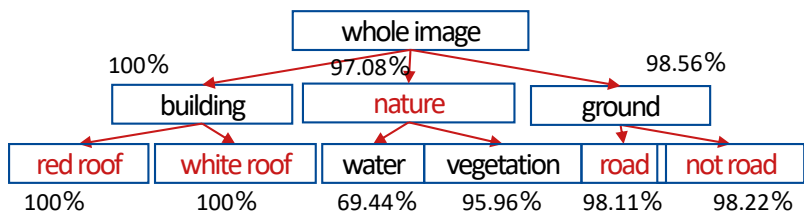
Figure 9: Segmentation results of virtual participant 2 on histopathological image.

original image’s five labels: building, clutter, vegetation, water and ground are removed and will not be presented to the participant. We collect 600 responses from the first iteration and 400 responses from the following 4 iterations.

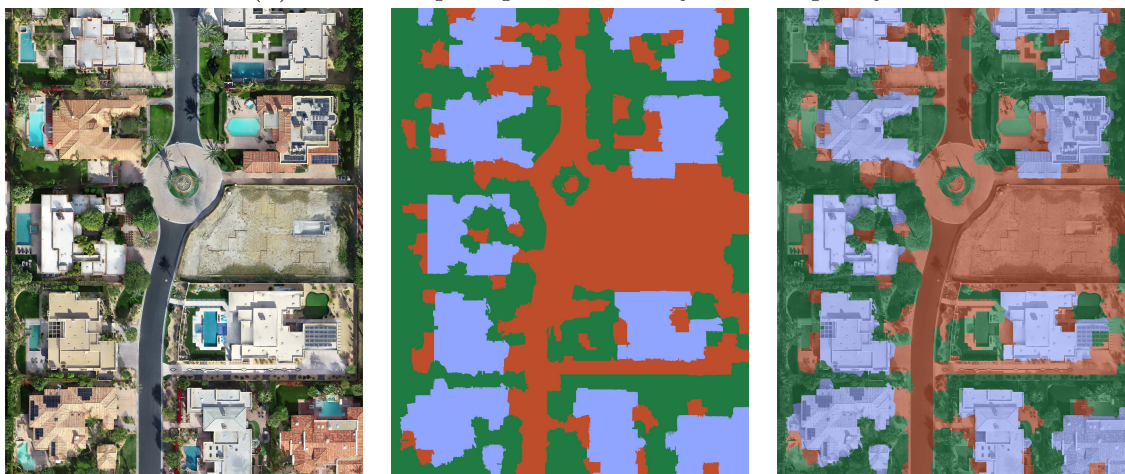
We present the results in Fig 10. A knowledge hierarchy is extracted based on the psychometric responses. We can notice that several concepts are discovered beyond the original labels, such as buildings divided into two sub-segmentation based on the colors of the roofs, the grounds split by whether vehicles could be driven on, water and vegetation merged together to natural habitats. We manually labeled each patch with these concepts and calculated the purity of each cluster. The purities are all above 95% except the water, because the area of water is too small to generate enough super-pixels for psychometric tests. Note that in the second layer segmentation, we could tell the concept relationships directly by the color similarities. For instance, in Fig 10(f)subfigure, the color of water and vegetation are all blue because they are all natural habitats, while vegetation is dark blue and water is light blue since they are different categories.

4.4. Comparison and Ablation Study

Two benchmarks are designed for comparison. First, we label each super-pixel by a set of fixed labels (the bottom layer’s labels in the hierarchy) for training and cluster on the model’s activation layer to build hierarchical segmentation. Another is the hierarchical segmentation proposed by [Arbelaez et al. \(2010\)](#). Dendrogram purity is applied to measure the quality of hierarchical segmentation against the ground truth knowledge. The results are reported in Table 1. It could be noticed that on both datasets, our method could get the best performance, indicate the capability of hierarchically segmenting based on the annotator’s high fidelity knowledge. It is no surprise that baseline 1 (hierarchical segmentation ([Arbelaez et al., 2010](#))) performs worst because no semantic perception is captured from the annotator. The results of baseline 2 (fixed label annotation) are much better but still worse than our framework since lacking the ability to extract full-depth knowledge. The gap between



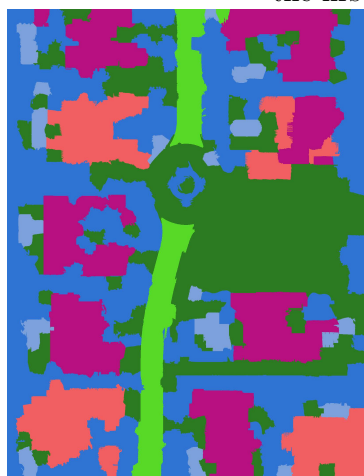
(a) Extracted perception hierarchy and the purity



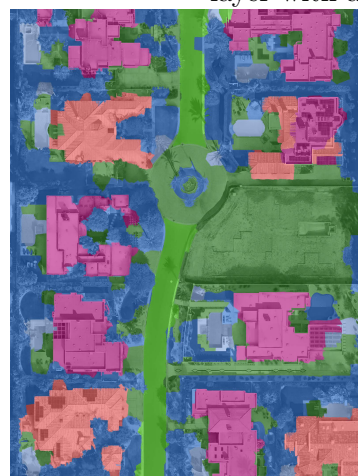
(b) Original image

(c) Visualize segmentation at the first layer

(d) Segmentation at the first layer with a mask



(e) Visualize segmentation at the second layer



(f) Segmentation at the second layer with a mask

Figure 10: Segmentation results of the real participant on the aerial image. (a) is the elicited knowledge structure by our method. Some new concepts are discovered beyond the original labels and are marked as red color. In (f), we could visually inspect concept relationships by colors. Close concepts are masked by similar colors and vice versa. For instance, red roof and white roof buildings are all masked red but with different color shades, because they are different categories but all belong to a bigger concept, building.

baseline 2 and the proposed method in histopathological image is smaller than that in synthetic image, as the hierarchical perception structure is more complex in the synthetic image, allowing our method to capture greater semantic knowledge.

Table 1: Comparison with baselines by dendrogram purity (%).

Participant	Dataset	Baseline 1	Baseline 2	Proposed method
1	Synthetic image	48.86	71.16	99.83
2	Synthetic image	44.71	73.19	93.47
1	Histopathological image	47.70	89.14	94.26
2	Histopathological image	48.63	88.81	97.97

To study the value of query selection and query enhancement strategies, we perform an ablation study on the synthetic image. We use the same experiment configuration as section 4.1. The training curve is shown in Fig 11. We can see that with the same training data, our results get a noticeable boost from random selection to active queries selection and a slight improvement with the triplet enhancement strategy, indicating the effectiveness of our proposed query scheme.

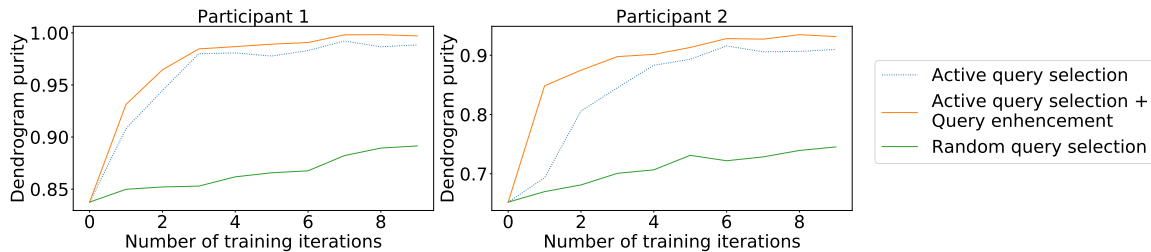


Figure 11: Ablation study on synthetic image.

5. Conclusion

In this work, we presented a method that captures the fine-grained semantic relationships between concepts in image segmentation using psychometric testing and deep metric learning. We empirically evaluated the ability of our method to effectively capture and represent the high fidelity semantics in synthetic, aerial and histology images. We present the results with a color overlay visualization where the captured concepts’ distances are reflected by the color distances in a given palette, which allows for an effective visual inspection. The results indicate the potential of this method to improve the elicitation of knowledge for image segmentation in a broad range of image analysis applications in domains such as biology, medicine, as industrial applications.

References

Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE transactions on pattern analysis and machine intelligence*, 34(11):2274–2282, 2012.

- Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010.
- Guilherme Aresta, Teresa Araújo, Scotty Kwok, Sai Saketh Chennamsetty, Mohammed Safwan, Varghese Alex, Bahram Marami, Marcel Prastawa, Monica Chan, Michael Donovan, et al. Bach: Grand challenge on breast cancer histology images. *Medical image analysis*, 56:122–139, 2019.
- Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- Weihua Chen, Xiaotang Chen, Jianguo Zhang, and Kaiqi Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 403–412, 2017.
- Lawrence T DeCarlo. On a signal detection approach to m-alternative forced choice with bias, with maximum likelihood and bayesian approaches to estimation. *Journal of Mathematical Psychology*, 56(3):196–207, 2012.
- Gustav Theodor Fechner. *Elemente der psychophysik*, volume 2. Breitkopf u. Härtel, 1860.
- Hsin-Chang Feng, Michael W Marcellin, and Ali Bilgin. A methodology for visually loss-less jpeg2000 compression of monochrome stereo images. *IEEE Transactions on Image Processing*, 24(2):560–572, 2014.
- Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3146–3154, 2019.
- George A Gescheider. *Psychophysics: the fundamentals*. Psychology Press, 2013.
- Ali Ghodsi. Dimensionality reduction a short tutorial. *Department of Statistics and Actuarial Science, Univ. of Waterloo, Ontario, Canada*, 37(38):2006, 2006.
- Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 1735–1742. IEEE, 2006.
- Katherine A Heller and Zoubin Ghahramani. Bayesian hierarchical clustering. In *Proceedings of the 22nd international conference on Machine learning*, pages 297–304, 2005.
- Matjaž Jogan and Alan A Stocker. A new two-alternative forced choice method for the unbiased characterization of perceptual bias and discriminability. *Journal of Vision*, 14(3):20–20, 2014.
- Mahmut Kaya and Hasan Şakir Bilge. Deep metric learning: A survey. *Symmetry*, 11(9):1066, 2019.

- Anna Khoreva, Rodrigo Benenson, Jan Hosang, Matthias Hein, and Bernt Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 876–885, 2017.
- Trupti M Kodinariya and Prashant R Makwana. Review on determining number of cluster in k-means clustering. *International Journal*, 1(6):90–95, 2013.
- Wei Liu, Andrew Rabinovich, and Alexander C Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- Lior Rokach and Oded Maimon. Clustering methods. In *Data mining and knowledge discovery handbook*, pages 321–352. Springer, 2005.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- Burr Settles. Active learning literature survey. 2009.
- Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1857–1865, 2016.
- IY Son, M Winslow, B Yazici, and XG Xu. X-ray imaging optimization using virtual phantoms and computerized observer modelling. *Physics in Medicine & Biology*, 51(17):4289, 2006.
- Yumin Suh, Bohyung Han, Wonsik Kim, and Kyoung Mu Lee. Stochastic class-based hard example mining for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7251–7259, 2019.
- Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7158–7166, 2017.
- Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020.
- Lu Yin, Vlado Menkovski, and Mykola Pechenizkiy. Knowledge elicitation using deep metric learning and psychometric testing. *arXiv preprint arXiv:2004.06353*, 2020.