



Active Learning for Segmentation by Optimizing Content Information for Maximal Entropy

Firat Ozdemir¹(✉), Zixuan Peng¹, Christine Tanner¹, Philipp Fuernstahl²,
and Orcun Goksel¹

¹ Computer-assisted Applications in Medicine, ETH Zurich, Zurich, Switzerland
ozdemirf@vision.ee.ethz.ch

² CARD, University Hospital Balgrist, University of Zurich, Zurich, Switzerland

Abstract. Segmentation is essential for medical image analysis tasks such as intervention planning, therapy guidance, diagnosis, treatment decisions. Deep learning is becoming increasingly prominent for segmentation, where the lack of annotations, however, often becomes the main limitation. Due to privacy concerns and ethical considerations, most medical datasets are created, curated, and allow access only locally. Furthermore, current deep learning methods are often suboptimal in translating anatomical knowledge between different medical imaging modalities. Active learning can be used to select an informed set of image samples to request for manual annotation, in order to best utilize the limited annotation time of clinical experts for optimal outcomes, which we focus on in this work. Our contributions herein are two fold: (1) we enforce domain-representativeness of selected samples using a proposed penalization scheme to maximize information at the network *abstraction layer*, and (2) we propose a Borda-count based sample querying scheme for selecting samples for segmentation. Comparative experiments with baseline approaches show that the samples queried with our proposed method, where both above contributions are combined, result in significantly improved segmentation performance for this active learning task.

1 Introduction

Segmentation has several medical applications, such as patient-specific surgical planning. Due to limited resources of expert physicians, detailed manual annotations are often not possible, even when desired anatomy may be visible with sufficient contrast using non-invasive imaging modalities such as MRI and ultrasound. Deep learning has shown encouraging performance for segmentation [1, 2], but often only when sufficient amount of labeled data for a target anatomy is available. Medical image data across different medical centers is often not uniform, for instance with respect to machine manufacturer, imaging settings, and cohort demographics. Thus, studies and corresponding annotations are only carried out in isolated datasets, with difficulties in merging information with data

sharing, patient rights, and confidentiality concerns. Hence, a sufficiently large dataset for a given task needs to be labeled. *Active learning* aims at maximizing the prediction performance through an intelligent sample querying system so that the limited expert annotation resources can be properly managed as opposed to training on a randomly selected next batch of samples which would contain a lot of redundancy. In a clinical environment, one can imagine that expert(s) will allocate a fixed amount of annotation time per time interval (i.e., week), hence the correct use of this time (i.e., on most valuable samples) is essential. Therefore, the segmentation framework would be initially provided a very limited labeled dataset, which will be extended with a certain batch size of samples intelligently selected at each iteration of the active learning.

Intuitively, the prediction confidence of a learned model can be used as a surrogate metric for its potential accuracy, in order to propose the most *uncertain* predictions for future manual annotation. In [3], *MC dropout* is proposed to sample from the approximate trained model posterior, which can be used to quantify an *uncertainty* metric through variations in the model predictions for a given input. Based on this, several approaches of querying the next batch of data are studied and compared with uniform random sampling in [4]. Unfortunately, it is intractable to assess conditional uncertainty of multiple samples; e.g. would i^{th} sample be still as uncertain as before once j^{th} sample is queried and trained for. Thus, it is intuitive to select a *representative* subset of these uncertain samples to reduce redundancy. Using a simplified version of DCAN [2] architecture (which has won the first place in the 2015 MICCAI Gland Segmentation Challenge [5]) for the purpose of faster training, a state-of-the-art method was proposed in [6] to select optimal sample images to annotate. First, a batch of *uncertain* samples is chosen based on the mean variance of multiple network predictions, followed by picking a subset of these using *maximum set coverage* [7] over the *image descriptors* of these samples. Recently in [8], a *content distance* [9] concept was proposed to quantify the similarity between two images, for selecting representative samples in class-incremental learning.

Herein we propose two main novelties for querying samples at an active learning step: (1) we add an additional constraint on the *abstraction layer* [8] activations during training to maximize information content at this level. We show that this additional constraint improves sample suitability that boosts segmentation performance from active learning. (2) Instead of the two step sample querying procedure (i.e., first select based on *uncertainty*, then cull using *representativeness*), we propose a Borda-count based method. This alone provides improvement over the state-of-the-art [6]; and when used in conjunction with our novel constraint above, it yields even further segmentation improvement.

2 Estimating Surrogate Metrics for Representativeness

Background. In [6], multiple FCNs were trained to estimate uncertainty for a given image through variation in their inferences. To make the FCN predictions diverse, the annotated dataset was also bootstrapped when training each model.

However, training several models is a costly operation and with larger number of models, one should bootstrap a smaller portion of the already-minimal dataset available in the early stages of typical active learning scenarios.

In our work, as a baseline, we implemented an improved version of the *Suggestive Annotation* framework [6]. We added dropout layers (c.f. Fig. 1) to allow for MC dropout [3], through which one can compute the voxel-wise variance across n_i inferences, and average it over all input voxels. The first step in querying samples is to pick the most uncertain n_{unc} samples S_{unc} from the set of non-annotated data D_{pool} . For representativeness, “image descriptor” I_i^c of every image $I_i \in D_{\text{pool}}$ is computed as described in [6] at the abstraction layer, l_{abst} (c.f. Fig. 1). Using cosine similarity $d_{\text{sim}}(I_i, I_j) = \cos(I_i^c, I_j^c)$ between the descriptors of images I_i and I_j , the maximum set-cover [7] over D_{pool} is computed using descriptors from S_{unc} for the top n_{rep} images. We call this method of using uncertainty and the above image descriptor (ID) as UNC-ID hereafter.




Fig. 1. DCAN network for *Suggestive Annotation* with additional spatial dropout layers. n_{ch} is the number of filters in respective block, BN is batch normalization, and n_{cl} is the number of classes. In consecutive bottlenecks, the first uses convolution filter in shortcuts to match tensor size while the second does not.

Content Distance. The image descriptor I_i^c averages the spatial information at the corresponding layer activations. While this allows for a spatially invariant means of representing a given image at a very abstract level, higher order features extracted at this stage would be blurred by this process. Alternatively, layer activation responses $R^l(I_i)$ of a pretrained classification network at a layer l can be used to describe the content of an image I_i [9]. Then, content distance (d_{cont}) between images I_i and I_j is defined as the mean squared error between their responses at layer l :

$$d_{\text{cont}}(I_i, I_j) = \frac{1}{N} \sum (R^l(I_i) - R^l(I_j))^2 \quad (1)$$

A similar notion can be applied to active learning problems, where input images are described by the activation response at the l_{abst} of the currently trained network (c.f. Fig. 1).

Encoding Representativeness by Maximizing Entropy. Content distance defined in Eq. (1) allows for finer content discrimination than image descriptors [6]. However, it has been suggested that activations at a single layer may not be sufficient for accurate content description [8]. This is likely to particularly apply to segmentation networks, since network weights until l_{abst} are not optimized to describe the input image. Therefore, it has been proposed to stack activations from multiple layers. For a typical segmentation network, storing all layer activations of D_{pool} can quickly diverge to an unfeasible size. Alternatively, one can try to increase information content at the l_{abst} through maximizing its activation entropy [10] along the feature channels. Entropy loss can then be defined as:

$$L_{\text{ent}} = - \sum_x \text{H}(R^{(l_{\text{abst}}, x)}) \quad (2)$$

where $R^{(l_{\text{abst}}, x)}$ are the input activations of all channels for spatial location x , and x iterates over the width and height of the layer l_{abst} . Hence, total loss for the trained network becomes $L_{\text{total}} = L_{\text{seg}} + \lambda L_{\text{ent}}$, where L_{seg} is the segmentation loss, and λ is used to scale the entropy loss L_{ent} .

Optimization of the network weights through entropy maximization is a novel regularization. L_{ent} alone would have a tendency to alter network weights to only increase information, which may also encourage randomness. With an appropriate λ , the network is forced to optimize parameters for the segmentation task while also increasing “useful” information content at the abstraction layer; as opposed to producing just noise at l_{abst} . Hence, additional content description for a given image can be retrieved from a single layer activation, making it a feasible alternative. We refer to this method, where an entropy-based content distance (**ECD**) is used, as **UNC-ECD**.

3 Sample Selection Strategy

For active learning, one should emphasize that the initial data size can be very small. Until the model parameters are optimized for a sufficient coverage of the data distribution, the defined “uncertainty” metric might be misleading. As a result, one can explore different ways to combine multiple metrics when querying samples instead of the conventional 2-step process. An intuitive way to combine two metrics m_k and m_l would be to use $w_k m_k + w_l m_l$, where w_k, w_l are weights. However, *uncertainty* and *representativeness* metrics defined in Sect. 2 are not linearly combinable, even if normalized, due to non-linear unit increments. Therefore, we propose to use Borda count, where samples are ranked for each metric, and the next query sample I_{i^*} is picked based on the best combined rank:

$$i^* = \arg \min_i \left(\sum_{m_k \in S_m} f_{\text{rank}}(m_k(I_i)) \right) \quad (3)$$

where S_m is the set of metrics m_k to combine, and the f_{rank} function sorts the images based on the metric m_k . When we use the ranking in Eq. (3) for samples

selection, we denote this in our results with “+”, e.g. content distance with uncertainty is named UNC+**E**CD. In an active learning framework, the methods mentioned until now can be denoted as UNC+ID, UNC+**E**CD for ranking based sample selection and UNC-ID, UNC-**E**CD for uncertainty selection followed by representativeness selection.

Table 1. Dataset configuration

Config	#volumes	Left/Right	vox res. [mm]	image size [px]	TR [s]	TE [s]	FA [°]
1	20	9/11	0.91 × 0.91 × 3.0	192 × 192 × 64	20	1.70	10
2	16	8/8	0.83 × 0.83 × 3.0	144 × 144 × 56	20	2.39	10

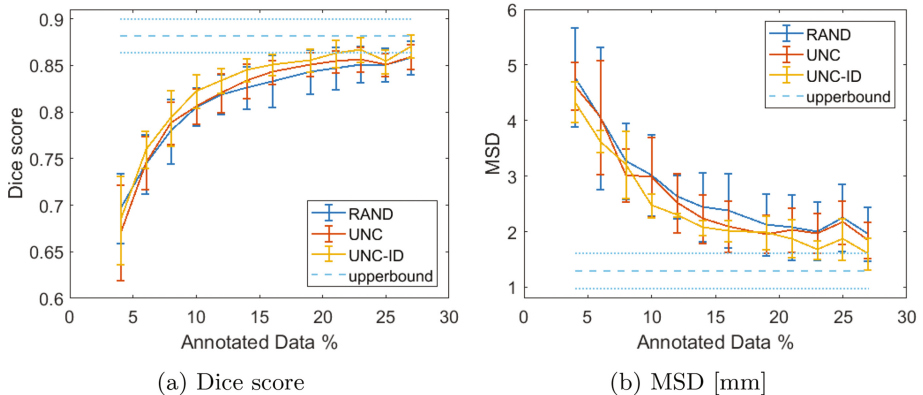


Fig. 2. Comparison between our implementation of the baseline method (UNC-ID) with random sampling (RAND) and only uncertainty-based (UNC) active learning methods. Training on 100% of the data (D_{pool}) is shown as upperbound. (a) Mean Dice score and (b) mean surface distance (MSD) with error bars covering the standard deviation of 5 hold-out experiments at every evaluation point.

4 Experiments and Results

We have conducted experiments on an MR dataset of 36 patients diagnosed with rotator cuff tear (shoulders) according to specifications shown on Table 1. In an effort to regularize the dataset, Config2 images have been resized to match the voxel resolution of Config1, and then zero padded to match the in-plane image size of Config1. The data has expert annotations of two bones (humerus & scapula) and two muscle groups (supraspinatus & infraspinatus + teres minor). Experiments have been conducted using NVIDIA Titan X GPU and Tensorflow library [11].




Fig. 3. Comparison of the baseline method (UNC-ID) with ranking based sample selection (UNC+ID) and the combination of our proposed extensions (UNC+ECD). Training on 100% of the data (D_{pool}) is shown as upperbound. (a) Mean Dice score and (b) mean surface distance (MSD) with error bars covering the standard deviation of 5 hold-out experiments at every evaluation point. The mean Dice score of UNC+ECD was statistically significantly higher than the baseline in 4 of 5 experiments (one-sided paired t-test at the 0.05 level).

For all compared methods, we have used the modified DCAN architecture shown in Fig. 1, training it on 2D in-plane slices with the parameters $n_{ch} = 32$ and Adam optimizer. When training the networks, learning rate of 5×10^{-4} , dropout rate of 0.5, $n_i = 17$, and minibatch size of 8 images were applied. At each active learning stage, including the initial training, models were trained for 8000 steps, which took about 48 mins. Uncertainty metric is aggregated over the foreground classes to represent their mean uncertainty. We used cross-entropy loss at the softmax layer (c.f. Fig. 1) for the L_{seg} . Weight λ for scaling L_{ent} in methods UNC-ECD and UNC+ECD is empirically set to $\lambda = 1/(360 \times |R^{\text{labst}}|)$.

To provide quantitative results, we have evaluated Dice score coefficient and mean surface distance (MSD). In an effort to efficiently utilize the available dataset, we have generated 5 hold-out experiments where the initial training set D_{an} , the non-annotated set D_{pool} , the validation set (all slices from 2 patients) and the test set (all slices from 9 patients) are randomly picked. All experiments are initially trained on 64 slices. For every active learning step, $n_{\text{rep}} = 32$ and $n_{\text{unc}} = 64$ is used. In Figs. 2 and 3, we show the Dice score and MSD of different methods evaluated for the test set at 11 stages of active learning ranging from 4% up to 27% of the D_{pool} . Conducted experiments are shown in two groups to increase clarity: (1) Comparison of our implementation of the baseline (UNC-ID) to uniform random sample querying (RAND) and sample querying based only on uncertainty (UNC) as seen in Fig. 2; (2) Building on top of (1), improvements of ranking (UNC+ID) and the gain from L_{ent} during training and representativeness capabilities of d_{cont} for sample querying, UNC+ECD (c.f. Fig. 3).




Fig. 4. Segmentation of a test volume comparing baseline (UNC-ID) with proposed method (UNC+ECD) after the first active learning step. Segmentation of two muscles overlaid on GS annotation (red) for (b) baseline and (c) proposed method. (d) Some of the substantial differences are pointed out by red arrows. (Color figure online)

In Fig. 4, we show an example cross-section from a test volume, where segmentation superiority of our proposed method (UNC+ECD) when compared to baseline is already visible after a single active learning step.

We conducted one-sided paired-sample t-tests at the 5% significance level on the mean Dice scores over all active learning steps for each hold-out experiment for UNC+ECD being superior to UNC-ID. Performance of UNC+ECD was statistically significantly better in 4 of 5 experiments.

5 Discussions and Conclusions

At early steps of active learning, one can see that the only uncertainty-based query sampling method (UNC) performs similar to random sample querying (RAND), with UNC only improving soon after $\approx 12\%$ of D_{pool} is used in training (c.f. Fig. 2). While UNC-ID already yields better segmentation performance than just uncertainty-based sampling, by simply using ranking, one can see that

the baseline method achieves a more substantial boost at early stages of active learning (see UNC+ID in Fig. 3). This behavior suggests that the surrogate uncertainty metric can give a bad approximation when the trained data size is fairly low; i.e., initial step(s). However, the suboptimal segmentation performance gain can be compensated with representativeness, and even further improved when given a higher priority; i.e., ranking instead of 2-step sample querying.

Upon combination of the proposed additional information maximization constraint during training and ranking combined with content distance at sample querying (UNC+ECD), we have observed the best Dice score on average at all active learning steps among the compared baseline and ranking extensions of the baseline methods. Other possible combinations of our proposed extensions (UNC-CD, UNC+CD, UNC-ECD) yielded inferior performance to UNC+ECD, and hence are not included in the quantitative comparisons to reduce clutter.

In this paper, we have comparatively studied the impact of different sample selection methods in active learning for segmentation. We have proposed 2 novel ways to query samples for active learning, which also can be combined to further boost performance during active learning steps. Compared to a state-of-the-art method, we have shown our proposed method to yield statistically significant improvement of segmentation Dice scores.

Acknowledgements. This work was funded by the Swiss National Science Foundation (SNSF), a Highly Specialized Medicine (HSM2) grant of the Canton of Zurich, and the EU's 7th Framework Program (Agreement No. 611889, TRANS-FUSIMO). We acknowledge NVIDIA GPU Grant support.

References

1. Baumgartner, C.F., Koch, L.M., Pollefeys, M., Konukoglu, E.: An Exploration of 2D and 3D Deep Learning Techniques for Cardiac MR Image Segmentation. In: Pop, M., et al. (eds.) STACOM 2017. LNCS, vol. 10663, pp. 111–119. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75541-0_12
2. Chen, H., Qi, X., Yu, L., Dou, Q., Qin, J., Heng, P.A.: DCAN: deep contour-aware networks for object instance segmentation from histology images. *Med. Image Anal.* **36**, 135–146 (2016)
3. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: International Conference on Machine Learning (ICML), pp. 1050–1059 (2016)
4. Gal, Y., Islam, R., Ghahramani, Z.: Deep Bayesian active learning with image data. In: Proceedings of the 34th International Conference on Machine Learning (ICML) (2017)
5. Sirinukunwattana, K., Pluim, J.P., Chen, H., et al.: Gland segmentation in colon histology images: the GlaS challenge contest. *Med. Imag Anal.* **35**, 489–502 (2017)
6. Yang, L., Zhang, Y., Chen, J., Zhang, S., Chen, D.Z.: Suggestive annotation: a deep active learning framework for biomedical image segmentation. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 399–407. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_46

7. Feige, U.: A threshold of $\ln n$ for approximating set cover. *ACM* **45**, 634–652 (1998)
8. Ozdemir, F., Fuernstahl, P., Goksel, O.: Learn the new, keep the old: extending pretrained models with new anatomy and images. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pp. 361–369. Springer International Publishing, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_42. ISBN: 978-3-030-00937-3
9. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *IEEE CVPR*, pp. 2414–2423 (2016)
10. Shannon, C.E.: A mathematical theory of communication. *ACM SIGMOBILE Mob. Comput. Commun. Rev.* **5**, 3–55 (2001)
11. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., et al.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). www.tensorflow.org