
Cost-Effective Active Learning for Melanoma Segmentation

Marc Gorriz
Xavier Giro-i-Nieto
Universitat Politecnica de Catalunya
Barcelona, Catalonia/Spain
xavier.giro@upc.edu

Axel Carlier
Emmanuel Faure
University of Toulouse
Toulouse, France
axel.carlier@enseeiht.fr

Abstract

We propose a novel Active Learning framework capable to train effectively a convolutional neural network for semantic segmentation of medical imaging, with a limited amount of training labeled data. Our contribution is a practical Cost-Effective Active Learning approach using dropout at test time as Monte Carlo sampling to model the pixel-wise uncertainty and to analyze the image information to improve the training performance.

1 Motivation

One of the major problems in medical diagnosis is the subjectivity of the specialist's decisions. More specifically, in the fields of medical imaging interpretation, the experience of the specialist can greatly determine the outcome of the final diagnosis. Manual inspection of medical images can in some cases be very tedious, time-consuming and subject to errors on part of the interpreter. This has led to a growth of computer vision algorithms to support medical diagnostics. The wide success of deep learning in computer vision has increased the accuracy of the predictions, renovating the interest in computer-assisted medical diagnosis.

Nevertheless, deep convolutional neural networks (CNNs) are defined by a huge amount of trainable parameters that require very large amounts of labeled data for training. This may be a heavy handicap in the medical imaging field, where the annotation costs of highly qualified professionals make very challenging the generation of large labeled datasets. Active Learning (AL) is an established approach to reduce this labeling workload in order to select, in an iterative way, a subset of informative examples from a large collections of unlabeled images. The chosen candidates are labeled and subsequently added to the training set.

There exist multiple active learning methods that choose which samples are to be annotated by the human experts. In this work we explore the uncertainty of the pixel-wise predictions as a selection criterion.

The main contributions of this work are the following:

- The design and training of a framework for medical imaging semantic segmentation using Convolutional Neuronal Networks (CNN) and the Cost-Effective Active Learning method.
- The development of information interpreters for medical imaging based on *Monte Carlo Dropout* for the analysis of the intrinsic network distribution.

The source code of this work is publicly available at <https://github.com/imatge-upc/medical-2017-nipsw>.

2 Related work

2.1 Cost-Effective Active Learning (CEAL) algorithm

An active learning is an algorithm able to interactively query the human annotator (or some other information source) new labeled instances from a pool of unlabeled data. Candidates to be labeled can be chosen with several methods based on informativeness and uncertainty of the data. Opposite to common Active Learning approaches that only consider the most informative and representative samples, a Cost-Effective methodology Wang et al. [2016] proposes to automatically select and pseudo-annotate unlabeled samples. A CEAL progressively feeds the samples from the unlabeled dataset into the CNN, and selects two kinds of samples for fine-tuning according to the CNN’s output. The sampling selection procedure is also known as complementary sample selection. One type of samples are a minority samples with low prediction confidence, which correspond to the most informative/uncertain samples whose prediction score is low. The selected samples are added into the labeled set after an active user (oracle) labels them. Another type of samples are the majority samples with high prediction score, because their prediction is highly confident. For these certain kind of samples, the proposed CEAL automatically assigns pseudo-labels with no human cost. These two types of samples are complementary to each other for representing different confidence levels of the current model on the unlabeled dataset.

2.2 CNN’s for Image Segmentation: U-Net architecture

Our active learning approach is applied over a popular convolutional neural network (CNN) in the medical domain: the U-Net. The U-Net Ronneberger et al. [2015] is a convolutional neural network architecture designed to solve biomedical image segmentation tasks. It was successfully used for winning the ISBI cell tracking challenge in 2015. The network combines a convolutional network architecture with a deconvolutional architecture to obtain the semantic segmentation Noh et al. [2015], Krähenbühl and Koltun [2011]. The convolutional network is composed of a repetitive pattern of two 3×3 convolutions operations, followed by a ReLU layer and a downsampling process through a 2×2 maxpooling operation with stride 2. On the other hand, the deconvolutional architecture includes a upsampling operation of the feature map obtained during the contracting path, followed by a 2×2 deconvolution that fractions the feature map channels into 2. A posteriori concatenation of the resulting feature map and the obtained during the contracting path is needed, followed by a 3×3 convolutions and a ReLU layer. The entire network is 23 convolutional layers deep, where the last layer is used to map each component feature vector related to the number of classes.

3 Proposed methodology

3.1 Image uncertainty estimation

The active learning criteria used to perform the complementary sample selection is based on the intrinsic distribution of the unlabeled data, ranking the unlabeled pool D^U based on their influence on the model. Being D_x^U, D_y^U the unlabeled data and their labels respectively, we are interested in finding the posterior network distribution $p(W|D_x^U, D_y^U)$. In general, this posterior distribution is not tractable, therefore we need to approximate the distribution of these weights using variational inference Graves [2011]. This technique allows us to learn the distribution over the network’s weights, $q(W)$, by minimizing the *Kullback-Leibler (KL)* divergence between this approximating distribution and the full posterior one.

In Kendall et al. [2015] it is explored the possibility to use *Monte Carlo Dropout* to approximate $q(W)$ term. The dropout works by randomly deactivating network activations following a *Bernoulli* distribution with a basis probability p_d per layer. With all the above, the introduction of *Dropout* during the training prevents overfitting Gal and Ghahramani [2015], while at test time it allows introducing pixel-wise sample uncertainty. Being I_x a image pixel, we can estimate the uncertainty of its predicted label \tilde{I}_y computing the variance of T different predictions on the same pixel by the effect of *Dropout* through the network weights. The precision of pixel-wise uncertainty maps will depend on the *Dropout* steps T and the *Dropout* probability p_d . High p_d means high variation of network weights, making difficult a consistent result with a finite number of step predictions. As




Figure 1: Pixel-wise uncertainty map using 10 step predictions.

was shown in Kendall et al. [2015] an optimal p_d value is 0.5 and a maximum precision would be obtained when $T \rightarrow \infty$.

In order to integrate the method to the CEAL complementary sample selection, the pixel-wise uncertainty must be transformed into a numerical score to estimate the prediction confidence. We propose summing up all the pixel values from the uncertainty map, obtaining this way higher scores for those most doubtful segmentations.

However, a simple pixel-wise addition does not cope with the variability of predictions across the contours. The further from the contour, the greater should be the contribution to the overall uncertainty of the prediction. For this reason, we introduce an additional weighting step based on the distance map over the predicted segmentation, which corresponds to computing the euclidean distance of each pixel to the closest pixel of the contour Breu et al. [1995]. This distance map is multiplied with the uncertainty map, so that the values of the further pixels from the predicted contour are boosted. Intuitively, this processing maker thicker the thickest contours in the uncertainty maps, and thinner the thinnest ones.

3.2 Complementary sample selection

Once the uncertainty score is defined, we can visualize its relation with the accuracy of the predictions. Figure 2 (left) shows the correlation between both, using as predictor the best model found in the next Section 4. The plot allows distinguishing four types of samples:

1. Undetected melanomas: although they have low uncertainty, they are also the most informative candidates to be manually annotated because they correspond to missing detections but with high certainty.
2. Highly uncertain samples: secondary candidates to be annotated by the oracle.
3. Certain samples: the most common ones, best candidates to be selected as a pseudo-labels.
4. Uncertain and wrong predictions: the most undesirable case whose population aims at be reduce by iterating the active learning algorithm.




Figure 2: Right: Regions uncertainty representation. Left: histogram projection.

Notice that in the real case scenario, the accuracy of the predictions cannot be computed because the ground truth of the analyzed samples is not available. In a real case only the uncertainty values can be estimated. A histogram of these uncertainty values is depicted in 2 (right), which corresponds to a projection of 2 (left) over the uncertainty axis. In this case, samples between the third and fourth regions become shuffled and cannot be discriminated. With all the above, the complementary data selection method will plan a strategy to select sequentially the correct proportion of samples for each region in each iteration, to improve gradually the network performance, without falling into over-fitting.

4 Results and Future work

All the tests were done with the ISIC 2017 Challenge dataset for Skin Lesion Analysis towards melanoma detection Gutman et al. [2016]. Each image in this dataset has a pixel-wise segmentation of the melanomas. In order to simulate an active learning scenario, our experiments only used a small portion of available annotations. The amount of segments used for training the CNN would increase based on the decisions of the proposed active learning approach, as if they were generated online by human annotators.

The dataset contains 2,000 RGB dermoscopy images manually annotated by medical experts, by tracing the lesion boundaries in the form of a binary mask (International Skin Imaging Collaboration). The dataset was modified for this work, transforming the original images to gray scale and modifying their aspect ratio to fit the CNN input dimensions. Before starting the interactive learning process, the training sets are initialized based on the Cost-Effective Active Learning methodology. First, it is randomly selected the labeled amount D^L from the dataset and then the other labels are deleted initializing the unlabeled set D^U .

The CNN was trained by applying data augmentation over the training samples as in Krizhevsky et al. [2012], to improve the generalization of the model.

The selection of samples in each iteration of the active learning loop was made according to the heuristic parameters summarized in the tables of Figure 3. The iterative algorithm starts by training a CNN with 600 samples and iteratively selects among a pool of 1,000 images which ones are to be labeled and added to the training set. In each iteration, the proposed algorithm samples 10 images where among those where no melanoma was detected, 10 more images among those with a highest uncertainty score, and 15 more images selected randomly. Also, those predictions with a confidence score over a certain threshold was selected as pseudo-labels and also included in the training set. Figure 3 also depicts the selected samples at some iterations 0, 4, 5 and 9.




Figure 3: Left: Regions evolution. Red samples showing the manual annotations. Right: Cost-Effective Active Learning data balance and heuristic parameters.

The quality of the segmentations was assessed based on the Dice Coefficient as in Ronneberger et al. [2015], which is presented in Equation 1:

$$\text{DICE}(A, B) = \frac{2|A\hat{B}|}{|A| + |B|} \quad (1)$$

where A represents the predicted masks for image and B its ground truth mask.

The results reached a Dice Coefficient Ronneberger et al. [2015] of 74 % after 9 active learning iterations with 2 training epochs over the CNN per run. Results indicate that after these 9 iterations there were still samples in the central region (4) that may limit the performance of the model.

Acknowledgements

This research was supported by contract SGR1421 by the Catalan AGAUR office. The work has been developed in the framework of project TEC2016-75976-R, funded by the Spanish Ministerio de Economía y Competitividad and the European Regional Development Fund (ERDF). The authors also thank NVIDIA for generous hardware donations.

References

- Heinz Breu, Joseph Gil, David Kirkpatrick, and Michael Werman. Linear time euclidean distance transform algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):529–533, 1995.
- Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158*, 2015.
- Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pages 2348–2356, 2011.
- David Gutman, Noel CF Codella, Emre Celebi, Brian Helba, Michael Marchetti, Nabin Mishra, and Allan Halpern. Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1605.01397*, 2016.
- Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv preprint arXiv:1511.02680*, 2015.
- Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Advances in neural information processing systems*, pages 109–117, 2011.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1520–1528, 2015.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 2016.