



PHiSeg: Capturing Uncertainty in Medical Image Segmentation

Christian F. Baumgartner¹(✉), Kerem C. Tezcan¹, Krishna Chaitanya¹,
Andreas M. Hötker², Urs J. Muehlemaier², Khoschy Schawkat^{2,4},
Anton S. Becker^{2,3}, Olivio Donati², and Ender Konukoglu¹

¹ Computer Vision Lab, ETH Zürich, Zürich, Switzerland
baumgartner@vision.ee.ethz.ch

² University Hospital Zürich, Zürich, Switzerland

³ Memorial Sloan Kettering Cancer Center, New York, USA

⁴ Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, USA

Abstract. Segmentation of anatomical structures and pathologies is inherently ambiguous. For instance, structure borders may not be clearly visible or different experts may have different styles of annotating. The majority of current state-of-the-art methods do not account for such ambiguities but rather learn a single mapping from image to segmentation. In this work, we propose a novel method to model the conditional probability distribution of the segmentations given an input image. We derive a hierarchical probabilistic model, in which separate latent variables are responsible for modelling the segmentation at different resolutions. Inference in this model can be efficiently performed using the variational autoencoder framework. We show that our proposed method can be used to generate significantly more realistic and diverse segmentation samples compared to recent related work, both, when trained with annotations from a single or multiple annotators. The code for this paper is freely available at <https://github.com/baumgach/PHiSeg-code>.

1 Introduction

Semantic segmentation of anatomical structures and pathologies is a crucial step in clinical diagnosis and many downstream tasks. The majority of recent automated segmentation methods treat the problem as a one-to-one mapping from image to output mask (e.g. [6]). However, medical segmentation problems are often characterised by ambiguities and multiple hypotheses may be plausible [10]. This is in part due to inherent uncertainties such as poor contrast or other restrictions imposed by the image acquisition, but also due to variations in annotation “styles” between different experts. To account for such ambiguities it is crucial

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-32245-8_14) contains supplementary material, which is available to authorized users.




Fig. 1. (Left) Example of hierarchical segmentation generation with segmentation output at each level (\hat{s}_ℓ) in the top row, right to left, and corresponding residual refinements to the prostate peripheral zone class in the bottom row. (Right) Corresponding graphical model (for N independent samples).

that prediction systems provide access to the full distribution of plausible outcomes without sacrificing accuracy. Predicting only the most likely hypothesis may lead to misdiagnosis and may negatively affect downstream tasks.

Recent work proposed to account for the uncertainty in the learned model parameters using an approximate Bayesian inference over the network weights [2]. However, it was shown that this method may produce samples that vary pixel by pixel and thus may not capture complex correlation structures in the distribution of segmentations [4]. A different line of work accounts for the possibility of different outcomes by training an ensemble of M networks [5] or by training a single network with M heads [7]. Both approaches, however, can only produce a fixed number of hypotheses. This problem is overcome by the conditional variational autoencoder (cVAE), an extension of [3] for modelling conditional segmentation masks given an input image [8]. Finally, the recently proposed probabilistic U-NET combines the cVAE framework with a U-NET architecture [4]. The authors showed that, given ground-truth annotations from multiple experts, the method can produce an unlimited number of realistic segmentation samples. Moreover, the method was shown to outperform various related methods including network ensembles, M -heads [7] and the Bayesian SegNet [2].

However, as we will show, the probabilistic U-NET produces samples with limited diversity. We believe this may be due to the fact that stochasticity is only introduced in the highest resolution level of the U-NET, and because the network can choose to ignore the random draws from the latent space since it is only concatenated to the channels. In this work, we propose a novel hierarchical probabilistic model which can produce segmentation samples closely matching the ground-truth distribution of a number of annotators. Inspired by Laplacian Pyramids, the model generates image-conditional segmentation samples by generating the output at a low resolution and then continuously refining the distribution of segmentations at increasingly higher resolutions. In contrast to prior work, the variations on *each resolution level* are governed by a separate latent variable, thereby avoiding the problems mentioned above. This process is illustrated in Fig. 1. We show that compared to recent work, our proposed Probabilistic Hierarchical Segmentation (PHiSeg) produces samples of significantly better quality for two challenging segmentation tasks, both, when trained with

multiple annotations, and a single annotation per image. Furthermore, the mean prediction of our model performs on par with the standard U-NET in terms of segmentation accuracy.

2 Methods

We start by assuming that the segmentations s given an input image x are generated from L levels of latent variables z_ℓ according to the graphical model shown in Fig. 1. Thus, the conditional distribution $p(s|x)$ is given by the following expression for the general case of L latent levels:

$$p(\mathbf{s}|\mathbf{x}) = \int p(\mathbf{s}|\mathbf{z}_1, \dots, \mathbf{z}_L)p(\mathbf{z}_1|\mathbf{z}_2, \mathbf{x}) \cdots p(\mathbf{z}_{L-1}|\mathbf{z}_L, \mathbf{x})p(\mathbf{z}_L|\mathbf{x})d\mathbf{z}_1 \cdots d\mathbf{z}_L. \quad (1)$$

We further assume that each latent variable z_ℓ is responsible for modelling the conditional target segmentations at $2^{-\ell+1}$ of the original image resolution (e.g. \mathbf{z}_1 and \mathbf{z}_3 model the segmentation at the original and at 1/4 of the original resolution, respectively.). This does not result from the graphical model itself but is rather enforced by our implementation thereof as will become clear shortly.

We aim to approximate the posterior distribution of $p(\mathbf{z}|\mathbf{s}, \mathbf{x})$ using a variational approximation $q(\mathbf{z}|\mathbf{s}, \mathbf{x})$ where we used \mathbf{z} to denote $\{\mathbf{z}_1, \dots, \mathbf{z}_L\}$. It can be shown that $\log p(\mathbf{s}|\mathbf{x}) = \mathcal{L}(\mathbf{s}|\mathbf{x}) + \text{KL}(q(\mathbf{z}|\mathbf{s}, \mathbf{x})||p(\mathbf{z}|\mathbf{s}, \mathbf{x}))$, where \mathcal{L} denotes the *evidence lower bound*, and $\text{KL}(\cdot, \cdot)$ the Kullback-Leibler divergence [3, 4, 8]. Since $\text{KL}(\cdot, \cdot) \geq 0$, \mathcal{L} is a lower bound on the conditional log probability with equality when the approximation q matches the posterior exactly. Using the decomposition in Eq. 1 we find that for our model

$$\begin{aligned} \mathcal{L} = & \mathbb{E}_{q(\mathbf{z}_1, \dots, \mathbf{z}_L|\mathbf{x}, \mathbf{s})} [\log p(\mathbf{s}|\mathbf{z}_1, \dots, \mathbf{z}_L)] - \alpha_L \text{KL} [q(\mathbf{z}_L|\mathbf{s}, \mathbf{x})||p(\mathbf{z}_L|\mathbf{x})] \\ & - \sum_{\ell=1}^{L-1} \alpha_\ell \mathbb{E}_{q(\mathbf{z}_{\ell+1}|\mathbf{s}, \mathbf{x})} [\text{KL} [q(\mathbf{z}_\ell|\mathbf{z}_{\ell+1}, \mathbf{s}, \mathbf{x})||p(\mathbf{z}_\ell|\mathbf{z}_{\ell+1}, \mathbf{x})]], \end{aligned} \quad (2)$$

with $\alpha_\ell = 1$. A complete derivation can be found in Appendix A. The α_ℓ are additional heuristic variables which we introduced to help account for dimensionality differences between the \mathbf{z}_ℓ (explained below). Following standard practice we parametrise the prior and posterior distributions as axis aligned normal distributions $\mathcal{N}(\mathbf{z}|\mu, \sigma)$. Specifically, we define

$$p(\mathbf{z}_\ell|\mathbf{z}_{\ell+1}, \mathbf{x}) = \mathcal{N} \left(\mathbf{z}|\phi_\ell^{(\mu)}(\mathbf{z}_{\ell+1}, \mathbf{x}), \phi_\ell^{(\sigma)}(\mathbf{z}_{\ell+1}, \mathbf{x}) \right) \quad (3)$$

$$q(\mathbf{z}_\ell|\mathbf{z}_{\ell+1}, \mathbf{x}, \mathbf{s}) = \mathcal{N} \left(\mathbf{z}|\theta_\ell^{(\mu)}(\mathbf{z}_{\ell+1}, \mathbf{s}, \mathbf{x}), \theta_\ell^{(\sigma)}(\mathbf{z}_{\ell+1}, \mathbf{s}, \mathbf{x}) \right), \quad (4)$$

where the ϕ, θ are functions parametrised by neural networks. Note that in contrast to the variational autoencoder [3], the $p(\mathbf{z}_\ell|\cdot, \mathbf{x})$ are also parametrised by neural networks similar to [4, 8]. Lastly, we model $p(\mathbf{s}|\mathbf{z})$ as the usual categorical distribution with parameters (i.e. softmax probabilities) predicted by another

neural network. By parametrising all distributions using neural networks, this can be seen as a hierarchical conditional variational autoencoder with the posteriors $q(\mathbf{z}_\ell|\cdot, \mathbf{s}, \mathbf{x})$ and priors $p(\mathbf{z}_\ell|\cdot, \mathbf{s})$ encoding \mathbf{x} and \mathbf{s} into latent representations \mathbf{z}_ℓ , and the likelihood $p(\mathbf{s}|\mathbf{z})$ acting as the decoder. Our implementation of this model using a neural network for $L = 3$ is shown in Fig. 2. In that figure it can be seen that the total number of resolution levels of the network (i.e. number of downsampling steps plus one) can be larger than the number of latent levels. The example in Fig. 2 has a total of 4 resolution levels, of which only $L = 3$ are latent levels. We obtained the best results with 7 total resolution levels of which $L = 5$ are latent levels. The prior and posterior nets have identical structure but do not share any weights. Similar to previous work, all three subnetworks are used for training but testing is performed by using only the prior and the likelihood networks [3, 4, 8].

From Fig. 2 it can be seen that latent variables \mathbf{z}_ℓ will form the skip connections in a U-NET-like architecture. However, unlike [6] and [4], each skip connection corresponds to a latent variable \mathbf{z}_ℓ such that no information can flow from the image to the segmentation output without passing a sampling step. We do not map the latent variables to a 1-D vector but rather choose to keep the structured relationship between the variables. We found that this substantially improves segmentation accuracy. As a result, latent variable \mathbf{z}_ℓ has a dimensionality of $r_x 2^{-\ell+1} \times r_y 2^{-\ell+1} \times D$, where D is a hyper-parameter and $D = 2$ for all experiments, and r_x, r_y are the dimensions of the input images. The latent variable \mathbf{z}_ℓ is limited to modelling the data at $2^{-\ell+1}$ of the original resolution due to the downsampling operations before it. It then passes up the learned representation to the latent space embedding above ($\mathbf{z}_{\ell-1}$) to perform a refinement at double the resolution. This continues until the top level is reached. To further enforce this behaviour the likelihood network is designed to generate only *residual* changes of the segmentation masks for all \mathbf{z}_ℓ except the bottom one. This is achieved through the addition layers before the outputs (see Fig. 2). Our model bears some resemblance to the Ladder Network [9] which is also a hierarchical latent variable model where inference results in an autoencoder with skip connections. Our work differs substantially from that work in how inference is performed. Furthermore, to our knowledge, the Ladder Network was never applied to structured prediction problems.

Training and Predictions: We aim to find the neural network parameters which maximise the lower bound \mathcal{L} in Eq. 2. The analytical form of the individual terms is prescribed by our model assumptions: since the posterior and prior both are modelled by normal distributions, the KL terms can be calculated analytically [3]. Our choice of likelihood results in a cross entropy term $\text{CE}(\hat{\mathbf{s}}_1, \mathbf{s}_{gt})$, with $\hat{\mathbf{s}}_1$ the predicted segmentation and \mathbf{s}_{gt} the corresponding ground-truth. Similar to previous work we found that it is sufficient to evaluate all of the expectations using a single sample [3]. Two deviations from the above theory were necessary for stable training. First, the magnitude of the KL terms depends on the dimensionality of \mathbf{z}_ℓ . However, since the dimensionality of \mathbf{z}_ℓ in our model grows with $O(2^\ell)$, this led to optimisation problems. To counteract this, we heuristically set




Fig. 2. Schematic network architecture of the proposed method for $L = 3$ latent levels and 4 resolution levels.

the weights $\alpha_\ell = 2^{\ell-1}$ in Eq. 2. Secondly, to enforce the desired behaviour that \mathbf{z}_ℓ should only model the data at its corresponding resolution, we added deep supervision to the output of each resolution level (\hat{s}_ℓ in Fig. 2). The cost function used for this is again the cross entropy loss, $\text{CE}(\text{ups}(\hat{s}_\ell), \mathbf{s}_{gt})$ for $\ell > 1$, where $\text{ups}(\cdot)$ denotes a nearest neighbour upsampling to match the size of \mathbf{s}_{gt} . While \mathbf{z}_ℓ can only model the data at a certain resolution, it may ignore this responsibility and focus only on matching the prior and posterior. Deep supervision effectively prevents this behaviour.

We trained the model using the Adam optimiser with a learning rate of 10^{-3} and a batch-size of 12. We used batch-normalisation on all non-output layers. All models were trained for 48 h on a NVIDIA Titan Xp GPU and the model with the lowest total loss on a held-out validation set was selected.

After the model is trained, segmentation samples for an input image \mathbf{x} can be generated by first obtaining samples \mathbf{z}_ℓ using the prior network and then decoding them using the likelihood network.

3 Experiments and Results

We evaluated our method on two datasets: (1) the publicly available LIDC-IDRI dataset which comprises 1018 thoracic CT images with lesions annotated by 4 radiologists [1]. Similar to [4] we extracted square 2D patches of size 128×128 pixels such that each patch was centred on a lesion. (2) We also evaluated our method on an in-house prostate MR dataset of 68 patients acquired with a transverse T2-weighted sequence (in-plane resolution $0.1875 \times 0.1875 \text{ mm}^2$ and slice thickness 3.3 mm). The transition and peripheral zones were manually annotated by 4 radiologists and 2 non-radiologists. We processed the data slice-by-slice (approx. 25 slices per volume), where we resampled each slice to a resolution of $0.6 \times 0.6 \text{ mm}^2$ and took a central crop of size 192×192 . We divided both datasets into a training, testing and validation set using a random 60-20-20 split.

For all experiments we compared our method (PHiSeg) with $L = 5$ latent levels and a total of 7 resolution levels to the probabilistic U-NET [4]. In order to exclude network capacity as an explanation for performance differences, we aimed to model our network components as closely as possible after the probabilistic U-NET. We used batch normalisation layers for both methods which deviates from [4] but did not affect the results negatively. Furthermore, to demonstrate that modelling the segmentation problem at multiple resolution levels is beneficial, we also compared against a variation of PHiSeg with only $L = 1$ latent levels (i.e. no skip connections or latent space hierarchy). Lastly, for some experiments we compared to a deterministic U-NET using the same architecture as for the probabilistic U-NET but with no stochastic components.

We evaluated the techniques in two experiments. First, we trained the methods using the masks from all available annotators, where in each batch we randomly sampled one annotation per image. We were interested in assessing how closely the distribution of generated samples matched the distribution of ground-truth annotations. To this end, we used the generalised energy distance $D_{\text{GED}}^2(p_{gt}, p_s) = 2 \mathbb{E}[d(\mathbf{s}, \mathbf{y})] - \mathbb{E}[d(\mathbf{s}, \mathbf{s}')] - \mathbb{E}[d(\mathbf{y}, \mathbf{y}')]$, where d is 1 minus the intersection over union, i.e. $d(\cdot, \cdot) = 1 - \text{IoU}(\cdot, \cdot)$, and $\mathbf{s}, \mathbf{s}', \mathbf{y}, \mathbf{y}'$ are samples from the learned distribution p_s , and ground-truth distribution p_{gt} [4]. The GED reduces the sample quality to a single, easy-to-understand number but, as a consequence, cannot be interpreted visually. Therefore, we additionally aimed to produce pixel-wise maps showing variability among the segmentation samples. We found the expected cross entropy between the mean segmentation mask and the samples to be a good measure, i.e. $\gamma(s_i) = \mathbb{E}[\text{CE}(\bar{s}_i, s_i)]$ with i the pixel position and \bar{s}_i the mean prediction. γ is statistically similar to variance with the L2-distance replaced by CE. However, we believe it is more suitable for measuring segmentation variability. Examples of our γ -maps along with sample segmentations are shown in Fig. 3. We quantify how well the γ -maps for each method predict regions with large uncertainty using the average normalised cross correlation (NCC) between the γ -maps and the CE error maps obtained with respect to each annotator:

$$\mathcal{S}_{\text{NCC}}(p_{gt}, p_s) = \mathbb{E}_{\mathbf{y} \sim p_{gt}} [\text{NCC}(\mathbb{E}_{\mathbf{s} \sim p_s} [\text{CE}(\bar{\mathbf{s}}, \mathbf{s})], \mathbb{E}_{\mathbf{s} \sim p_s} [\text{CE}(\mathbf{y}, \mathbf{s})])]. \quad (5)$$




Fig. 3. Ground-truth annotations and samples for two of the evaluated methods trained with masks from 6 or 1 expert(s). Average error maps $\mathbb{E}_{y,s}[\text{CE}(y, s)]$ and γ -maps ($\mathbb{E}[\text{CE}(\bar{s}, s)]$) for each model are shown in the left and right-most column, respectively.

Results for both D_{GED}^2 and \mathcal{S}_{NCC} are shown in the top part of Table 1. All measures were evaluated with 100 samples drawn from the learned models.

Secondly, we set out to investigate the models’ ability to infer the inherent uncertainties in the annotations from *just one* annotation per training image. To this end, we trained the above models by using only the annotations of a single expert. For the evaluation we then computed the D_{GED}^2 and \mathcal{S}_{NCC} using all available annotators. Additionally, we evaluated the models in terms of conventional Dice score evaluated with masks from the single annotator as ground-truth. To get a single prediction from the probabilistic models we used \bar{s} . This allowed us to obtain an indication of conventional segmentation accuracy. The results are shown in the bottom part of Table 1.

We observed that when using all annotators for training, PHiSeg ($L = 5$) produced significantly better D_{GED}^2 and \mathcal{S}_{NCC} scores compared to all other methods. This can be observed qualitatively in Fig. 3 for a prostate slice with large inter-expert disagreements. Both, the prob. U-NET and PHiSeg ($L = 5$) produced realistic samples but PHiSeg ($L = 5$) was able to capture a wider vari-

Table 1. Quantitative results for all metric. Statistically significant improvements ($p < 0.01$ with paired student’s t -test) over all other methods are indicated in bold.

	# experts	LIDC-IDRI			Prostate dataset		
		D_{GED}^2	\mathcal{S}_{NCC}	Dice	D_{GED}^2	\mathcal{S}_{NCC}	Dice
Prob. U-NET	All	0.2393	0.7749	–	0.1322	0.7763	–
PHiSeg ($L = 1$)	All	0.2934	0.7944	–	0.1608	0.7452	–
PHiSeg ($L = 5$)	All	0.2248	0.8453	–	0.0864	0.8185	–
Det. U-NET	1	–	–	0.5297	–	–	0.8364
Prob. U-NET	1	0.4452	0.5999	0.5238	0.2198	0.6022	0.8290
PHiSeg ($L = 1$)	1	0.4695	0.6013	0.5275	0.2462	0.6683	0.7942
PHiSeg ($L = 5$)	1	0.3225	0.7337	0.5408	0.2044	0.6917	0.8540

ability. Furthermore, as indicated by the high \mathcal{S}_{NCC} values, PHiSeg’s ($L = 5$) γ -maps were found to be very predictive of where in the image the method’s average prediction errors will occur. Similar results were obtained when training with only one annotator. We noticed that in this scenario the prob. U-NET may in some cases fail to learn variation in the data and revert back to an almost entirely deterministic behaviour (see fourth row in Fig. 3). We believe this can be explained by the prob. U-NET’s architecture which, in contrast to our method, allows the encoder-decoder structure to bypass the stochasticity. While our method also predicted smaller variations in the samples, they were still markedly more diverse. The lower performance of PHiSeg ($L = 1$) indicates that using multiple resolution levels is crucial for our method. More samples for the prostate and LIDC-IDRI datasets can be found in Appendix B. From Table 1 it can be seen that no significant differences between the Dice scores were found for any of the methods (except PHiSeg’s ($L = 1$)), including the det. U-NET. From this we conclude that neither PHiSeg ($L = 5$) nor the prob. U-NET suffer in segmentation performance due to their stochastic elements.

4 Discussion and Conclusion

We introduced a novel hierarchical probabilistic method for modelling the conditional distribution of segmentation masks given an input image. We have shown that our method substantially outperforms the state-of-the-art on a number of metrics. Furthermore, we demonstrated that PHiSeg was able to predict its own errors significantly better compared to previous work. We believe that proper modelling of uncertainty is indispensable for clinical acceptance of deep neural networks and that having access to the segmentation’s probability distribution will have applications in numerous downstream tasks.

Acknowledgements. This work was partially supported by the Swiss Data Science Center. One of the Titan X Pascal used for this research was donated by the NVIDIA Corporation.

References

1. Armato, S.G., et al.: The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. *Med. Phys.* **38**(2), 915–931 (2011)
2. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian SegNet: model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. [arXiv:1511.02680](https://arxiv.org/abs/1511.02680) (2015)
3. Kingma, D., Welling, M.: Auto-encoding variational bayes. [arXiv:1312.6114](https://arxiv.org/abs/1312.6114) (2013)
4. Kohl, S., et al.: A probabilistic U-Net for segmentation of ambiguous images. In: *Proceedings of NIPS*, pp. 6965–6975 (2018)
5. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *Proceedings of NIPS*, pp. 6402–6413 (2017)
6. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
7. Rupprecht, C., et al.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In: *Proceedings of CVPR*, pp. 3591–3600 (2017)
8. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: *Proceedings of NIPS*, pp. 3483–3491 (2015)
9. Valpola, H.: From neural PCA to deep unsupervised learning. In: *Perspectives of Neural Computing*, pp. 143–171. Elsevier (2015)
10. Warfield, S.K., Zou, K.H., Wells, W.M.: Validation of image segmentation and expert quality with an expectation-maximization algorithm. In: Dohi, T., Kikinis, R. (eds.) *MICCAI 2002*. LNCS, vol. 2488, pp. 298–306. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45786-0_37