

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2019 ; 2019: 356–364. doi:10.1007/978-3-030-32251-9_39.

As Easy as 1, 2... 4? Uncertainty in Counting Tasks for Medical Imaging

Zach Eaton-Rosen^{1,2}, Thomas Varsavsky^{1,2}, Sebastien Ourselin², M. Jorge Cardoso²

¹Centre for Medical Imaging Computing, University College London, London, UK

²Biomedical Engineering and Imaging Sciences, King's College London, London, UK

Abstract

Counting is a fundamental task in biomedical imaging and count is an important biomarker in a number of conditions. Estimating the uncertainty in the measurement is thus vital to making definite, informed conclusions. In this paper, we first compare a range of existing methods to perform counting in medical imaging and suggest ways of deriving predictive intervals from these. We then propose and test a method for calculating intervals as an output of a multi-task network. These predictive intervals are optimised to be as narrow as possible, while also enclosing a desired percentage of the data. We demonstrate the effectiveness of this technique on histopathological cell counting and white matter hyperintensity counting. Finally, we offer insight into other areas where this technique may apply.

1 Introduction

Counting is a common analysis task required in a wide range of medical imaging applications from histology (cell counting) to neuroradiology (lesion counting). For any of these clinical biomarkers, accurate quantification of the degree of uncertainty over the measurements is of high importance in deciding an appropriate course of action. In this paper, we demonstrate an improved method for quantifying the uncertainty for counting tasks.

Uncertainty can be broadly broken down into two constituent parts: model and data uncertainty. In the context of CNNs, 'model' or 'epistemic' uncertainty represents the uncertainty over the network (weights, hyperparameters, architecture) while 'data' or 'aleatoric' uncertainty represents the noise inherently associated with the data (noisy labels, measurement noise). Furthermore, out-of-distribution examples are likely to adversely affect the performance of machine-learning tools.

Epistemic uncertainty can be assessed through the comparison of several samples obtained from stochastic neural networks. If the stochasticity is induced by dropout, the sampling

Correspondence to: Zach Eaton-Rosen.

zach.eaton-rosen@kcl.ac.uk.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-32251-9_39) contains supplementary material, which is available to authorized users.

approximates full Bayesian inference [1], which has been employed in image segmentation applications [2]. When training deep learning models, the stochasticity inherent in minibatching makes it possible to compare different models trained on the same training dataset: differing predictions of these models can be attributed to model uncertainty. A network can be trained to output diverse predictions from m ‘heads’ [3] coming from a common network trunk: the m heads’ differences are, again, due to model uncertainty.

Heteroscedastic models of the noise uncertainty have been used in image super-resolution [4] and also exploited for spatially adaptive task loss weighting in multi-task learning [5]. However, these parametric methods are restricted to unimodal, symmetric distributions, which are not necessarily realistic. Testtime augmentation has been used to perturb the data and thus infer the uncertainty from the differences in predictions [6,7]. In these approaches, the estimated uncertainty will depend wholly on the model’s lack of invariance to the chosen augmentations: this may suggest that the models are undertrained or lacking capacity.

While decomposing errors may be useful, other solutions exist. **Predictive Intervals** (PIs) estimate a lower and upper bound for an observation, such that the observation falls inside these bounds some chosen (high) percentage of the time: for a 95% PI, we would expect 95 of 100 observations to lie within the interval. PIs should satisfy the following properties: (a) to be as small as possible, while (b) still enclosing the appropriate fraction of results. This can be enforced through the loss function [8]. In this work, we propose an extension of this method for application to counting tasks. We first describe different methods to perform counting tasks and assess uncertainty over measurements before presenting the loss function. We then describe the proposed amendment which is more flexible and stabler to train.

2 Methods

2.1 Uncertainty in Counting

The overall aim of this work is to compute a predictive interval that (a) minimises the interval width, whilst (b) ensuring that the interval contains the appropriate percentage of results. Here we introduce several techniques to count cells and associated methods to calculate predictive intervals. While they present a novel contribution in their own right, these methods of counting are introduced here as a baseline against the method proposed in Sect. 2.3. These baseline methods do not explicitly regress a predictive interval. Instead, we use the multiple outputs from the models (e.g. MC samples or M -heads) to sample predictions of count. Although we could simply use percentiles of these results to calculate intervals, these perform badly (the true count is often not within the obtained bounds). In order to mitigate this issue and introduce a fair comparison, the predictive intervals for each method below are calibrated post hoc following [9]: we transform the bounds affinely until they encompass a fraction $f > 1 - \alpha$ on the validation data. This transformation is then applied to the test-set estimates.

Segmentation-Based—One counting method is to learn a segmentation of the input image and use connected-components analysis to determine the number of individual objects. To calculate uncertainty, we use three different approaches. We use Monte-Carlo samples of a network trained with dropout to produce N segmentation maps, counting the

objects for each of the N . Secondly, we measure the number of objects at different thresholds of the output confidence map of the network. As the confidence threshold increases, fewer pixels will remain in the ‘foreground’ class. Finally, we also use M -heads [3], which produces M estimates of the segmentation with one forward pass of the network. This method produces a diverse mode-seeking ensemble, so higher variability in the heads may indicate that the model is more uncertain of the segmentation. For these methods, our target is the segmentation from Fig. 1.

Regressing the Euclidean Distance Transform—Naylor et al. [10] introduced a regression-based method for cell counting. Given an input image, the network learns to approximate the Euclidean Distance Transform of the cell segmentation maps (see Fig. 1). A non-maximum suppression is then used to count the cells. In order to calculate the uncertainty of the cell count, we use the MC-sampling paradigm described above, with a network with dropout trained.

Regressing the Pixel-Wise Cell Density—One popular technique for counting in computer-vision applications is to use a regression formulation to estimate a density-map from the raw image: summing over all pixels returns the count [11,12]. The density estimation function we use is a convolutional neural network. For these experiments, the ground truth density map has value 0 for background, and $\frac{1}{n_i}$ otherwise, where n_i is the number of pixels in the i^{th} object. The network we use to estimate the density is a multi-task network with a shared backbone and two ‘heads’. One head learns the segmentation while the other learns the density map. For uncertainty estimation, we use the M -heads paradigm to introduce variance in the output.

2.2 Distribution-Free Uncertainty Estimation

The previous examples all rely on sampling, followed by post-training calibration step which maps the sample uncertainty to the target predictive bounds. Conversely, the authors in [8] proposed to regress the predictive intervals directly from the data. They aim to estimate lower and upper confidence bounds for a desired quantity y , where b_l and b_u are the lower and upper bounds respectively. The hyper-parameter a determines the desired width of the interval: it is defined such that:

- (i) $p(y \in [b_l, b_u]) = 1 - a$, and
- (ii) $p(y \in \cup\{(b_u, \infty], [-\infty, b_l)\}) = a$.

Common choices for a include 0.10, 0.05 and 0.01, representing 90%, 95% and 99% confidence intervals respectively. In the original work, the authors propose a loss function L_{QD} to estimate ‘Quality-Driven’, distribution-free predictive intervals. For any given datapoint, x_i , the model returns b_{l_i} and b_{u_i} . For each input X_i we assess whether the observed corresponding datapoint y_i is in or out of the prediction bounds $[b_{l_i}, b_{u_i}]$. In order to provide useful information on the behaviour of the predictive interval estimates over multiple examples, the loss function is allowed to reason over an entire minibatch of size n . In this setup, an indicator variable is used to express if y_i within the predictive interval or not. The number of times y_i falls within the predictive interval is given by a binomial distribution

Binomial($n, (1 - \alpha)$), assuming i.i.d. data, which can be approximated by a normal distribution for large n . The loss, L_{QD} , is then defined as the sum of a width term and the log-likelihood term.

$$\mathcal{L}_{QD} = \overline{W}_{captured} + \lambda \frac{n}{\alpha(1-\alpha)} \max(0, (1 - \alpha) - q)^2 \quad (1)$$

where $\overline{W}_{captured}$ is the mean interval width for intervals that capture their associated ground truth, λ is a constant, n is the number in a batch, and q is the fraction of points that lie within their estimated predictive interval bounds.

2.3 Proposed Extension

In practice, we found L_{QD} difficult to optimise. We observed periodic instabilities in the training and attributed this, in part, to the one-sided nature of the second term; we sought to modify its formulation appropriately. With this in mind, let P_s be a discrete probability distribution function representing the probability of being ‘in’ or ‘out’ of the predicted interval, where the subscript s denotes ‘state’, i.e. $P_{in} = 1 - \alpha$, $P_{under} = \alpha/2$ and $P_{over} = \alpha/2$. For any minibatch, we define the observed proportions of ‘over’, ‘under’ and ‘out’ samples as Q_s , and use the Kullback-Leibler divergence (KL) to enforce similarities between the target P and the observed Q . With this framework, Q could be encouraged to match any desired distribution P (for instance, estimating several bounds to correspond to percentiles). Note that, in this proposed framework, and contrary to [8], P can represent any chosen distribution.

$$\mathcal{L}_{distribution} = KL(P \parallel Q) = \sum_s P_s \log \frac{P_s}{Q_s} = \sum_s \{P_s \log(P_s) - P_s \log(Q_s)\} \quad (2)$$

Since the target distribution P is constant, minimizing $\mathcal{L}_{distribution}$ is identical to minimising the cross-entropy term with respect to the network weights; this means that we are simply promoting that the proportion of inliers in a given minibatch matches our desired distribution. As this loss function uses the categorical membership of the y_j to estimate Q , a soft membership function is used to make this operation differentiable and hence suitable for back-propagation.

We calculate the proportion inside the bounds as $Q_{in,soft} = \sigma(\xi(\mathbf{y} - \mathbf{b}_l)) \odot \sigma(\xi(\mathbf{b}_u - \mathbf{y}))$. The minibatch of ground truth counts \mathbf{y} is compared with the regressed bounds, \mathbf{b}_l and \mathbf{b}_u , with σ representing the sigmoid function and ξ being a positive softening constant (set to $\xi = 2$). This formulation of the soft boundaries is as in [8], with the other soft memberships (over, under) being set analogously.

Our proposed loss is given by:

$$\mathcal{L}_{proposed} = \overline{W}_{captured} + \lambda \sum_s P_s \log(Q_s) \quad (3)$$

In short, instead of using a one-sided data likelihood term, we used the crossentropy calculated between the chosen P and Q . This reformulation has not only added flexibility, allowing for different state chosen P distributions, but we also found it easier to train than the model in Eq. 1.

2.4 Network Architectures and Implementation Details

The U-Net [13] forms the basis of all of our CNNs. The multi-task network has a U-Net backbone with the same parameters as for the single-task approaches. It then splits into separate branches (one for segmentation and one for density prediction).

For the proposed method of uncertainty prediction, we fit a regression network with three output quantities. First, it outputs a ‘predicted count’ which is trained with an L_2 loss. The other two quantities are the upper and lower bounds, for a given a . In our experiments we choose $\alpha = 0.1$, making them 90% intervals. The architecture was chosen to have residual blocks as part of the arm to avoid vanishing gradients.

Due to the complexity of our network, we train it in stages. First, we train the U-Net part on the segmentation and density tasks. We then freeze the weights and train the predictive intervals, with a batch size of 64: as discussed, a large batch size to estimate good batch-wise statistics. In our experiments, we set $\lambda = 30$. The auxiliary L_2 loss is set to $1e-3$. We parametrise the outputs of the network as such: the estimate for the mean value has no final activation. The upper and lower residuals go through a softplus activation function and are then added or subtracted, as appropriate, to the mean estimate. The segmentation has a sigmoid output, and the density a square function. Models are trained with early stopping as determined on the validation set, in NiftyNet [14].

2.5 Data

The proposed counting methods are applied to the counting of cells from histological slides [10]. This dataset has 33 labelled slides of dimension 512×512 , taken from 7 different types of tissue, and each slide has an associated cell label map used here as the count ground truth. We separated this into 7 ‘test’ images (21%), one from each cell type, and 4 ‘validation’ images. To have larger batch sizes, we trained on images of size 256×256 and hence quartered each image while keeping the same fold label. Heavy augmentation is applied to the images in the training set (see figures in Supplementary Materials) using the ‘imgaug’ library [15].

We also fit to a white-matter hyperintensity (WMH) dataset [16]. In this task, we demonstrate a slightly different parameterisation of our bound prediction. We fitted an M -Heads model to the WMH segmentation and used the same model as a feature extractor to train the predictive bounds. In this data, of the 60 subjects, we used an 80/10/10 split for training/testing/validation respectively.

3 Results

We show the results in Table 1. All of the models exhibit good performance in counting, with the correlation between predicted and GT counts being above 0.8 for all models. The

uncalibrated predictive intervals capture anything from 9% to 61% of the data. After calibrating these models, many achieved the correct percentage of inliers for the predictive intervals. Some (for example, the M -heads density regression with $f_{in} = 0.75$) did not: this may indicate that the calibration methods were overfitting (despite having few parameters—only an affine transformation). Our model predicted significantly smaller interval widths than for the baseline methods for both cells and lesions. While the baseline methods may seem to give large bounds, in the cases of cells, there may be a count of over 100 per image and in the lesions, up to 35. Because the EDT regression had the lowest MAE, we chose another, simple, baseline: we simply had a percentage count as the error. This method achieved a width of 18.6, compared to 12.20 (ours) in Fig. 3. It also does not capture the desired percentage of inliers, as it is too small. For the model fitted with L_{QD} , we report the best results obtained after 3 independent model-fits, as we found the loss was unstable to fit—however, it still underperformed our proposed loss function.

4 Discussion

The aim of this work was to accurately predict intervals, such that they were of minimal width while containing the desired numbers of ground truth values. Our predictive bounds were over 20% smaller than the nearest competitor method while retaining the correct number of inliers; these smaller bounds are correspondingly more informative. We have demonstrated these results on a cell histology dataset and on WM lesions. For the cells, the next-best method was applying a constant percentage uncertainty to the counts of the EDT regression framework. Fitting this percentage is, in essence, optimising the same loss as we applied, but only using the predicted counts (and none of the image features). The fact that our model outperforms this baseline implies that the imaging features are being used to make an informed estimate of predictive error.

One limitation of this work is that it is not likely to generalise to samples drawn from outside of the training distribution. Domain-adaptation methods could help ameliorate this. It is also not an interpretable model and hence it would of interest to use model introspection methods to investigate how the network decides on its bounds. As the model we have presented can, in principle, be applied to any estimate derived from a machine-learning model, future work will investigate its applicability to 3D counting problems and a wider range of clinical biomarkers.

Acknowledgements

ZER is supported by the EPSRC Doctoral Prize. MJC & SO are supported by the Wellcome Flagship Programme (WT213038/Z/18/Z) and the Wellcome EPSRC CME (WT203148/Z/16/Z). We gratefully acknowledge NVIDIA Corporation for the donation of hardware.

References

1. Gal, Y; Ghahramani, Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning; International Conference on Machine Learning; 2016. 1050–1059.
2. Kendall A, Badrinarayanan V, Cipolla R. Bayesian SegNet: model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. arXiv preprint. 2015 arXiv:1511.02680

3. Lee S, Purushwalkam S, Cogswell M, Crandall D, Batra D. Why M heads are better than one: training a diverse ensemble of deep networks. arXiv preprint. 2015 arXiv:1511.06314
4. Tanno, R; , et al. In: Descoteaux, M; Maier-Hein, L; Franz, A; Jannin, P; Collins, DL; Duchesne, S, editors. Bayesian image quality transfer with CNNs: exploring uncertainty in dMRI super-resolution; MICCAI 2017. LNCS; 2017. 611–619.
5. Bragman, FJS; , et al. In: Frangi, AF; Schnabel, JA; Davatzikos, C; Alberola-López, C; Fichtinger, G, editors. Uncertainty in multitask learning: joint representations for probabilistic MR-only radiotherapy planning; MICCAI 2018. LNCS; 2018. 3–11.
6. Ayhan, MS; Berens, P. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks; MIDL; 2018.
7. Wang G, Li W, Aertsen M, Deprest J, Ourselin S, Vercauteren T. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. Neurocomputing. 2019; 21:34–45.
8. Pearce T, Zaki M, Brintrup A, Neely A. High-quality prediction intervals for deep learning: a distribution-free, ensembled approach. arXiv preprint . 2018 arXiv:1802.07167
9. Eaton-Rosen, Z; Bragman, F; Bisdas, S; Ourselin, S; Cardoso, MJ. In: Frangi, AF; Schnabel, JA; Davatzikos, C; Alberola-López, C; Fichtinger, G, editors. Towards safe deep learning: accurately quantifying biomarker uncertainty in neural network predictions; MICCAI 2018. LNCS; 2018. 691–699.
10. Naylor P, Laé M, Reyat F, Walter T. Segmentation of nuclei in histopathology images by deep regression of the distance map. IEEE Trans Med Imaging. 2019; 38(2):448–459. [PubMed: 30716022]
11. Lempitsky V, Zisserman A. Learning to count objects in images. Advances in neural information processing systems. 2010:1324–1332.
12. Xie W, Noble JA, Zisserman A. Microscopy cell counting and detection with fully convolutional regression networks. Comput Methods Biomech Biomed Eng Imaging Vis. 2018; 6(3):283–292.
13. Ronneberger, O; Fischer, P; Brox, T. In: Navab, N; Hornegger, J; Wells, WM; Frangi, AF, editors. U-Net: convolutional networks for biomedical image segmentation; MICCAI 2015. LNCS; 2015. 234–241.
14. Gibson E, et al. NiftyNet: a deep-learning platform for medical imaging. Comput Methods Programs Biomed. 2018; 158:113–122. [PubMed: 29544777]
15. Jung AB. imgaug. 2018
16. Kuijf H, et al. Standardized assessment of automatic segmentation of white matter hyperintensities; results of the WMH segmentation challenge. IEEE Trans Med Imaging. 2019 [PubMed: 30908194]

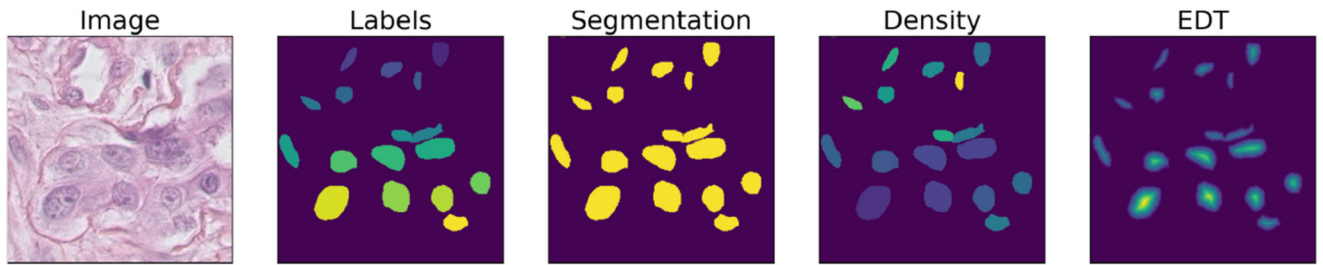
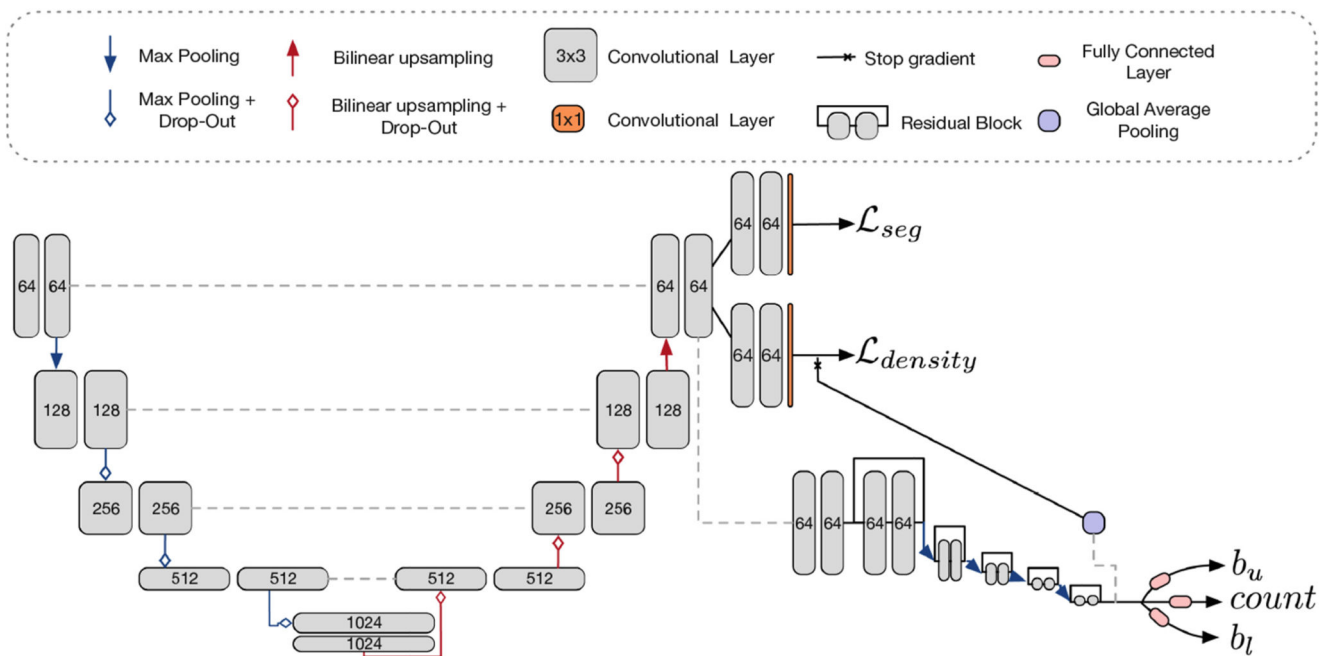


Fig. 1.

An illustrative figure of the cell data. The image has ground-truth labels that we binarise for the segmentation targets. The density and Euclidean Distance Transform (EDT) are targets for regression.

**Fig. 2.**

Multi-task architecture for simultaneous segmentation and uncertainty prediction. All convolutions are 3×3 by the channel width, denoted in the diagram. The U-Net is complemented by an ‘arm’ which has residual blocks followed by max-pooling (maintaining 64 filters) until it reaches the output layer, where it returns an upper and lower bound. Dropout is enabled for methods that require MC-sampling where indicated in the diagram, with $p = 0.5$. The bounds are trained with the loss from Eq. 2.

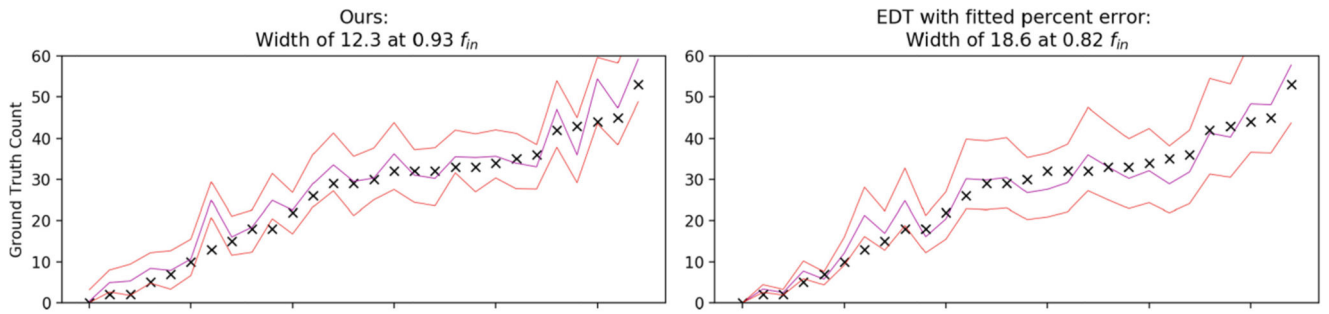


Fig. 3.

Here we contrast our model (left) with the model with fitted percentage noise. The points represent the ground truths, and the lines represent the upper bound, mean and lower bound (note that the lines are for ease of visualisation: the x -axis is not continuous).

Table 1

C_{EST} is the average estimated count. The ground truth counts were 25.71 for cells and 8.19 for lesions. MAE is the mean absolute error, \pm standard deviation. ρ is the correlation coefficient between estimated and ground truth counts. f is the fraction of the ground truth points within the bounds and W is the mean width of the intervals—evaluated on both uncalibrated and calibrated intervals.

Method	Paradigm	C_{EST}	MAE \pm STD	ρ	f_{uncal}	W_{uncal}	f_{cal}	W_{cal}
Segmentation	Thresholds	30.32	6.16 \pm 5.43	0.90	0.61	15.33	0.86	23.42
	M-Heads	25.14	2.83 \pm 3.34	0.95	0.46	2.80	0.96	102.22
	MC samples	31.36	6.70 \pm 6.37	0.87	0.43	7.25	0.89	27.82
EDT regression	MC samples	26.16	2.85 \pm 2.03	0.97	0.25	3.00	0.96	48.09
	% Errors	—	—	—	—	—	0.82	18.61
Density regression	M-Heads	26.42	3.53 \pm 3.12	0.96	0.57	6.42	0.75	15.64
L_{QD}	PI-estimate	26.01	3.04 \pm 2.82	0.96	—	—	1.0	29.31
Ours	PI-estimate	26.23	2.93 \pm 2.93	0.96	—	—	0.93	12.20
Lesions: segmentation	M-Heads	6.08	2.89 \pm 2.96	0.83	0.09	0.63	0.96	24.1
Ours	PI-estimate	6.08	2.89 \pm 2.96	0.83	—	—	0.89	10.93