



Assessing Reliability and Challenges of Uncertainty Estimations for Medical Image Segmentation

Alain Jungo^(✉)  and Mauricio Reyes

Insel Data Science Center, Inselspital, Bern University Hospital,
University of Bern, Bern, Switzerland
alain.jungo@artorg.unibe.ch

Abstract. Despite the recent improvements in overall accuracy, deep learning systems still exhibit low levels of robustness. Detecting possible failures is critical for a successful clinical integration of these systems, where each data point corresponds to an individual patient. Uncertainty measures are a promising direction to improve failure detection since they provide a measure of a system's confidence. Although many uncertainty estimation methods have been proposed for deep learning, little is known on their benefits and current challenges for medical image segmentation. Therefore, we report results of evaluating common voxel-wise uncertainty measures with respect to their reliability, and limitations on two medical image segmentation datasets. Results show that current uncertainty methods perform similarly and although they are well-calibrated at the dataset level, they tend to be miscalibrated at subject-level. Therefore, the reliability of uncertainty estimates is compromised, highlighting the importance of developing subject-wise uncertainty estimations. Additionally, among the benchmarked methods, we found auxiliary networks to be a valid alternative to common uncertainty methods since they can be applied to any previously trained segmentation model.

Keywords: Uncertainty · Segmentation · Deep learning

1 Introduction

Deep learning-based methods have led to impressive improvements in medical image segmentation over the past years. For many tasks, the performance is comparable to human-level performance, or even surpasses it [11]. Nonetheless, despite improvements in accuracy, the robustness aspects of these systems call for significant improvements for a successful clinical integration of these technologies, where each data point corresponds to an individual patient. This highlights

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-32245-8_6) contains supplementary material, which is available to authorized users.

the importance of having mechanisms to effectively monitor computer results in order to detect and react on system’s failures at the patient level. Among others, uncertainty measures are a promising direction since uncertainties can provide information as to how confident the system was on performing a given task on a given patient. This information in turn can be used to leverage the decision-making process of a user, as well as to enable time-effective corrections of computer results by for instance, focusing on areas of high uncertainty.

Different approaches have been proposed to quantify uncertainties in deep learning models. Among the most popular approaches are: (a) Bayesian uncertainty estimation via test-time dropout [5], (b) aleatoric uncertainty estimation via a second network output [9], and (c) uncertainty estimation via ensembling of networks [10]. In medical image segmentation, uncertainty measures are of interest at three levels. The first, most fine-grained level, is the voxel¹-wise uncertainty, which provides a measure of uncertainty for the predicted class of each voxel. This level of uncertainty is especially useful for the interaction with humans, be it by providing additional information to foster comprehensibility or as guidance for correction tasks. The second level is the uncertainty at the level of a segmented instance (or object). Nair et al. [12] and Graham et al. [6] used instance-level uncertainty to reduce the false discovery rate of brain lesions and cells, respectively. In both approaches voxel-wise uncertainties were aggregated to obtain an instance-wise uncertainty. Similarly, Eaton-Rosen et al. [4] aggregated voxel-wise uncertainties of brain tumor segmentations to obtain confidence intervals for tumor volumes. The third level is the subject-level uncertainty, which informs us whether the segmentation task was successful (e.g., above a certain metric). Having information about success or failure would be sufficient for many tasks, e.g., high-throughput analysis or selection of cases for expert review. As proposed by Jungo et al. [8], task-specific aggregation of the voxel-wise uncertainties could be used to obtain subject-level uncertainties. In contrast, DeVries et al. [3] and Robinson et al. [13], proposed an auxiliary neural network that predicts segmentation performance at the subject-level. A current challenge to use these latter type of approaches is that considerable large training datasets are necessary in practice to ensure their reliability [3].

In order to better understand the benefits and current challenges in uncertainty estimation for medical image segmentation, we evaluated common uncertainty measures with respect to their reliability, their benefit, and limitations². Additionally, we analyzed the requirements for uncertainties in medical image segmentation and we make practical recommendations for their evaluation.

2 Material and Methods

2.1 Data

We selected two publicly available, and distinct datasets for the experiments. The first dataset is the brain tumor segmentation (BraTS) challenge dataset

¹ For simplicity, we use *voxel* even if it could be a two-dimensional image.

² Code available at <https://github.com/alainjungo/reliability-challenges-uncertainty>.

2018 [1] consisting of 265 subjects. Each subject features four magnetic resonance images (T1-weighted, T1-weighted post-contrast, T2-weighted, FLAIR) of a size of $240 \times 240 \times 155$ isotropic (1 mm^3) voxels. We split the dataset into 100 training, 25 validation, and 160 testing subjects, combined the three tumor sub-compartment labels to segment the whole tumor, and performed a z-score intensity normalization ($\mu = 0, \sigma = 1$) on each subject and image individually. The second dataset is the international skin imaging collaboration (ISIC) lesion segmentation dataset 2017 [2] consisting of 2000 training, 150 validation, and 600 testing images. We resized the color images to a size of 256×192 pixels and normalized the intensities to the range $[0, 1]$.

2.2 Experimental Setting

Our aim is to evaluate the reliability of uncertainty measures for deep learning-based segmentation of medical images. Rather than building a specific fine-tuned, top-performing segmentation model, we used a U-Net-like architecture [14] due to its popularity, simplicity, and to minimize architectural influences on the outcomes³. The architecture consists of four pooling/upsampling steps and has dropout regularization ($p = 0.05$) and batch normalization after each convolution. We used a common training scheme consisting of a cross-entropy loss with Adam optimizer (learning rate: 10^{-4}), and applied early stopping with respect to the validation set Dice coefficient. Any adaptation to this architecture and training scheme was performed to fit the needs of each studied uncertainty approach.

2.3 Uncertainty Methods

We evaluated the following five different uncertainty measures:

Baseline Uncertainty: Softmax Entropy. Although the softmax output of a model is arguably a probability measure [5], we considered it as reference comparison as it is implicitly generated by segmentation networks. We named this strategy *baseline*. We used the normalized entropy $H = -\sum_{c \in \mathcal{C}} p_c \log(p_c) / \log(|\mathcal{C}|) \in [0, 1]$ as a measure of uncertainty, where p_c is the softmax output for class c and \mathcal{C} is the set of classes ($\mathcal{C} = \{0, 1\}$ in our case).

MC Dropout. Test time dropout can be viewed as an approximation of a Bayesian neural network [5]. T stochastic network samples can be interpreted as Monte-Carlo samples of the posterior distribution of the network’s weights and result in a class probability of $p_c = 1/T \sum_{t=1}^T p_{t,c}$. We employed the normalized entropy of these probabilities as a measure of uncertainty. For the experiments, we used $T = 20$ and considered two different dropout layer positioning strategies. First, we applied MC dropout on the base model (see Sect. 2.2), which uses minimal dropout ($p = 0.05$) after each convolution. Second, we applied more

³ We also conducted experiments with a DenseNet-like architecture with no notable differences in the outcome and therefore omit it here for space and clarity reasons.

prominent dropout ($p = 0.5$) at the center positions (i.e., before pooling and after upsampling, similar to [12]). Accordingly, we name these two strategies as *baseline+MC* and *center+MC*.

Aleatoric Uncertainty. In contrast to the model uncertainty (captured by e.g. MC dropout), the aleatoric uncertainty is said to capture noise inherent in the observation [9]. It is obtained by defining a network f with two outputs $[\hat{x}, \sigma^2] = f(x)$ and input x , where the outputs \hat{x} and σ^2 are the mean and variance of the logits perturbed with Gaussian noise. The aleatoric loss optimizes both outputs simultaneously by MC sampling (ten samples in our case) of the perturbed logits. We used \hat{x} for the class predictions and the variance σ^2 as a measure of uncertainty. We normalized the variance to $[0, 1]$ over all predictions.

Ensembles. Another way of quantifying uncertainties is by ensembling multiple models [10]. We combined the class probabilities of each network k by the average $p_c = 1/K \sum_{k=1}^K p_{k,c}$ over all $K = 10$ networks and used the normalized entropy as uncertainty measure. The individual networks share the same architecture (see Sect. 2.2) but were trained on different subsets (90%) of the training dataset and different random initialization to enforce variability.

Auxiliary Network. Inspired by [3, 13], where an auxiliary network is used to predict segmentation performance at the subject-level, we apply an auxiliary network to predict voxel-wise uncertainties of the segmentation model by learning from the segmentation errors (i.e., false positives and false negatives). For the experiments, we considered two opposing types of auxiliary networks. The first one, named *auxiliary feat.*, consists of three consecutive 1×1 convolution layers cascaded after the last feature maps of the segmentation network. The second auxiliary network, named *auxiliary segm.*, is a completely independent network (same U-Net as described in Sect. 2.2) that uses as input the original images and the segmentation masks produced by the segmentation model (generated by five-fold cross-validation). We normalized the output uncertainty subject-wise to $[0, 1]$ for comparability purposes.

2.4 Assessing Quality of Uncertainties

We adopted three metrics to evaluate the quality of uncertainties. Additionally, we computed the Dice coefficient to also verify segmentation performance as uncertainty methods typically link both tasks.

Calibration. Model calibration is important when not only the predicted class but also its corresponding confidence is of interest. In this regards, calibration has been used as a surrogate to assess the reliability of uncertainties [9]. A model is said to be perfectly calibrated if its predictions $f(x)$ with confidence p do occur with a fraction p of the time ($P(y = 1|f(x) = p) = p$ for the binary case). Meaning for example that for 100 predictions with a confidence of 0.7, 70 predictions are expected to be correct [7]. We assessed calibration of uncertainties by reliability diagrams and expected calibration error (ECE) [7]. Reliability diagrams show the deviation of the perfect calibration by plotting the binned

predicted confidences against the accuracy obtained for each bin (fraction of positives). The ECE is defined as the absolute error of these bins (i.e., the gap between confidence and accuracy) weighted by the number of samples in the bins, where a lower ECE (close to zero) indicates a better calibration. In our experiments, we used a bin size of ten and used the model output probabilities as confidence. For methods not providing segmentation probabilities but direct segmentation uncertainty estimates (i.e., auxiliary and aleatoric), we translated the uncertainties by $y(1 - 0.5q) + (1 - y)(0.5q)$ to confidences, where $y \in \{0, 1\}$ is the segmentation label and $q \in [0, 1]$ is the normalized uncertainty.

Uncertainty-Error Overlap. In a practical setting, perfect calibration of a model is impossible [7]. Often, segmentation tasks do not require perfect calibration but it would be sufficient for a model to be uncertain where it makes mistakes and certain where it is correct. To assess this condition, we used the overlap (determined by the Dice coefficient) between the segmentation error and the thresholded uncertainty, termed *uncertainty-error overlap* ($U-E$). This metric is not influenced by the true negatives from background areas, which are typically enormous in medical image segmentation. It is therefore an alternative for the ECE, which includes foreground as well as background areas.

Corrections. Motivated by previous works using uncertainty estimations, we assessed the quality of uncertainties by evaluating their benefit to correct segmentations. We define TPU, TNU, FPU, FNU as uncertainty in the true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). A beneficial correction is said to improve the Dice coefficient, hence, to benefit from removal of false positives, the relation $FPU(TP) > TPU(TP + FP + FN)$ needs to be satisfied (for the accuracy $FPU > TPU$ is sufficient). Similarly, in order to benefit from adding voxels (i.e., correct false negatives), the relation $FNU(TP + FP + FN) > TNU(TP)$, needs to be satisfied. However, the latter relation is not practically applicable due to large backgrounds and thus typically large TNU . Since voxel-wise corrections (as opposed to instance-wise corrections) might be more harmful than beneficial, we calculated the proportion of subjects that fulfill the benefit condition for false positive removal, BnF , as means of comparison to other methods.

3 Results

Figure 1 compares the calibration at the dataset level (i.e., all voxels in the dataset) with the calibration at the subject level (i.e., voxels of one subject). It shows the miscalibration that can occur at subject level (S1 and S2) while the calibration at dataset level is good. We found approximately 28%/46% underconfident and 32%/18% overconfident calibrations for the subjects of the BraTS/ISIC dataset. This underlines the special caution needed when using the calibration-based metrics (e.g., ECE) at the dataset level, as it can lead to misperception on the actual calibration quality of a model, and hence, the reliability of

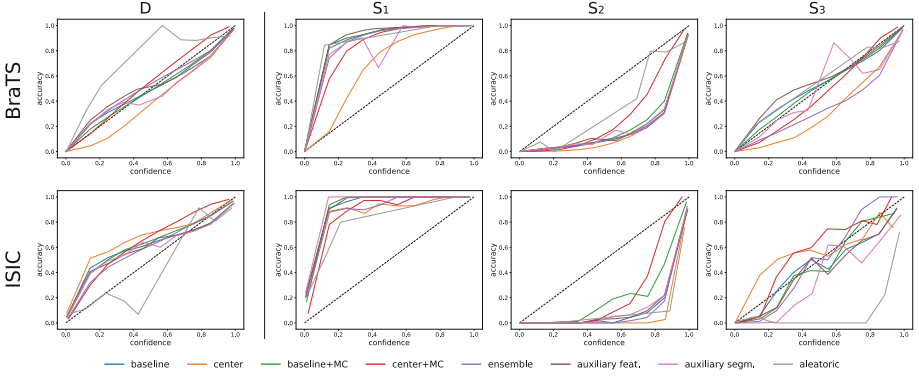


Fig. 1. Calibration at the dataset level (D) compared to the (mis)calibration at the subject level (S1, S2, S3) for the different uncertainty methods. S1, S2, S3 correspond to exemplary subjects for which the models are underconfident (S1), overconfident (S2), and well-calibrated (S3). Rows correspond to results on the BraTS and ISIC datasets.

its uncertainty estimations. Noticeable is also the agreement among the uncertainty methods at subject-level, suggesting only little benefit in selecting one uncertainty method over another.

In Table 1, we report for BraTS and the ISIC dataset the following metrics: average subject-level ECE (dataset-level ECE in supplementary material A), uncertainty-error overlap (U-E), proportion of correction-benefiting test subjects (BnF), and Dice coefficient. For a fair comparison, we selected the best-performing threshold for each method whenever the metric required an uncertainty threshold (i.e., U-E and BnF). Overall for both datasets, no uncertainty method outperforms and stands out over the others. Particularly, the *aleatoric* method and methods with large dropout (*center/+MC*) yield worst performance. The *aleatoric* method fails to produce uncertainty at the locations of segmentation errors (i.e., low U-E) and is therefore unable to improve segmentation results through corrections, whereas the large dropout mainly negatively affects segmentation performance and ECE. The results further show that MC dropout (*baseline+MC* and *center+MC*) typically improves ECE, U-E, and Dice coefficient over the non-MC versions (*baseline* and *center*), but larger amounts of dropout (*baseline < center* and *baseline+MC < center+MC*) results in worse performances, which suggests using MC dropout in the regimes where the benefit with respect to the uncertainty is minimal compared to standard softmax. We could confirm this finding through intermediate dropout strategies (see supplementary material B). We also observe good performances of the auxiliary networks, which are typically well-calibrated and profit from a good segmentation performance of their segmentation network (i.e., *baseline* model). In regards to the metrics, we note that low ECE values stem from large amount of low-confident background areas that positively affects the ECE. This also explains the lower ECE values for the BraTS dataset, which contains more background (even with applied brain mask) than the ISIC dataset, due to the additional

image dimension. Additionally, the BnF only considers TPU and FPU uncertainties and is therefore favorable for methods with low precision (more FP typically yields more FPU). We found this to be the reason for the bad correction performance of the ensemble on the BraTS dataset, even though the uncertainty-error overlap was good.

Table 1. Performances of the different uncertainties with respect to expected calibration error (ECE), uncertainty-error overlap (U-E), proportion of correction-benefiting test subjects (BnF), and Dice coefficient. Values are presented as *mean (rank)*. Standard deviation is omitted due to marginal differences. Upwards and downwards arrow indicate desired higher and lower metric values, respectively. Horizontal separation group types of uncertainty methods.

	BraTS				ISIC			
	ECE % ↓	U-E ↑	BnF ↑	Dice ↑	ECE % ↓	U-E ↑	BnF ↑	Dice ↑
Baseline	0.925 (4)	0.432 (2)	0.39 (3)	0.874 (2)	7.256 (4)	0.424 (4)	0.26 (4)	0.814 (3)
Center	1.758 (7)	0.409 (5)	0.5 (1)	0.866 (5)	9.415 (8)	0.411 (6)	0.27 (3)	0.78 (6)
Baseline+MC	0.9 (1)	0.433 (1)	0.36 (4)	0.874 (2)	7.36 (5)	0.428 (3)	0.24 (5)	0.813 (4)
Center+MC	1.233 (6)	0.433 (1)	0.27 (6)	0.868 (4)	8.766 (7)	0.428 (3)	0.17 (6)	0.794 (5)
Ensemble	0.919 (2)	0.433 (1)	0.32 (5)	0.879 (1)	7.131 (1)	0.431 (2)	0.31 (2)	0.831 (1)
Auxiliary feat	0.923 (3)	0.427 (3)	0.48 (2)	0.874 (2)	7.216 (3)	0.421 (5)	0.33 (1)	0.814 (3)
Auxiliary segm	0.925 (4)	0.412 (4)	0.48 (2)	0.874 (2)	7.212 (2)	0.433 (1)	0.27 (3)	0.814 (3)
Aleatoric	1.134 (5)	0.054 (6)	0.06 (7)	0.872 (3)	7.837 (6)	0.058 (7)	0.12 (7)	0.82 (2)

4 Discussion

The results show that although current voxel-wise uncertainty measures are rather well-calibrated at the dataset level (i.e., all voxels in the dataset) they tend to fail at the subject level (Fig. 1). This observation is to be expected since subject-level calibration errors (under- or overcalibration) can average out at the dataset level. Based on the proposed calibration-based metric, no overall best uncertainty measure was found among the studied methods. From our experiments we can conclude that methods that aggregate voxel-wise uncertainty to provide subject-level estimations are not reliable enough to be used as a mechanism to detect failed segmentations. We thus conclude on the importance of developing subject-level uncertainty estimation in medical image segmentation that can cope with the issue of High-Dimension-Low-Sample-Size (HDLSS) to ensure their reliability in practice.

Unsurprisingly, the ensemble method yields rank-wise the most reliable results (Table 1) and would typically be a good choice (if the resources allow it). The results also revealed that methods based on MC dropout are heavily dependent on the influence of dropout on the segmentation performance. In contrast, auxiliary networks turned out to be a promising alternative to existing uncertainty measures. They perform comparable to other methods but have the benefit

of being applicable to any high-performing segmentation network not optimized to predict reliable uncertainty estimates. No significant differences were found between using *auxiliary feat.* and *auxiliary segm.*. Through a sensitivity analysis performed over all studied uncertainty methods (see supplementary material C), we could confirm our observations that different uncertainty estimation methods yield different levels of precision and recall. Furthermore, we observed that when using current uncertainty methods for correcting segmentations, a maximum benefit can be attained when preferring a combination of low precision segmentation models and uncertainty-based false positive removal.

Our evaluation has several limitations worth mentioning. First, although the experiments were performed on two typical and distinctive datasets, they feature large structures to segment. The findings reported herein may differ for other datasets, especially if these consists of very small structures to be segmented. Second, the assessment of the uncertainty is influenced by the segmentation performance. Even though we succeeded in building similarly performing models, their differences cannot be fully decoupled and neglected when analyzing the uncertainty.

Overall, we aim with these results to point to the existing challenges for a reliable utilization of voxel-wise uncertainties in medical image segmentation, and foster the development of subject/patient-level uncertainty estimation approaches under the condition of HDLSS. We recommend that utilization of uncertainty methods ideally need to be coupled with an assessment of model calibration at the subject/patient-level. Proposed conditions, along with the threshold-free ECE metric can be adopted to test whether uncertainty estimations can be of benefit for a given task.

Acknowledgments. This work was supported by the Swiss National Foundation by grant number 169607. The authors thank Fabian Balsiger for the valuable discussions.

References

1. Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge. arXiv preprint [arXiv:1811.02629](https://arxiv.org/abs/1811.02629) (2018)
2. Codella, N.C., et al.: Skin lesion analysis toward melanoma detection: a challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In: ISBI, pp. 168–172. IEEE (2018)
3. DeVries, T., Taylor, G.W.: Leveraging uncertainty estimates for predicting segmentation quality. arXiv preprint [arXiv:1807.00502](https://arxiv.org/abs/1807.00502) (2018)
4. Eaton-Rosen, Z., Bragman, F., Bisdas, S., Ourselin, S., Cardoso, M.J.: Towards safe deep learning: accurately quantifying biomarker uncertainty in neural network predictions. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11070, pp. 691–699. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_78
5. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: representing model uncertainty in deep learning. In: ICML, pp. 1050–1059 (2016)

6. Graham, S., et al.: Mild-net: minimal information loss dilated network for gland instance segmentation in colon histology images. *Med. Image Anal.* **52**, 199–211 (2019)
7. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *ICML*, pp. 1321–1330. *JMLR. org* (2017)
8. Jungo, A., et al.: Uncertainty-driven sanity check: application to postoperative brain tumor cavity segmentation. *MIDL* (2018)
9. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: *NIPS*, pp. 5574–5584 (2017)
10. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *NIPS*, pp. 6402–6413 (2017)
11. Litjens, G., et al.: A survey on deep learning in medical image analysis. *Med. Image Anal.* **42**, 60–88 (2017)
12. Nair, T., Precup, D., Arnold, D.L., Arbel, T.: Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11070, pp. 655–663. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00928-1_74
13. Robinson, R., et al.: Real-time prediction of segmentation quality. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11073, pp. 578–585. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00937-3_66
14. 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