

Efficient Bayesian Uncertainty Estimation for nnU-Net

Zhao, Yidong; Yang, Changchun; Schweidtmann, Artur; Tao, Qian

DOI

[10.1007/978-3-031-16452-1_51](https://doi.org/10.1007/978-3-031-16452-1_51)

Publication date

2022

Document Version

Final published version

Published in

Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings

Citation (APA)

Zhao, Y., Yang, C., Schweidtmann, A., & Tao, Q. (2022). Efficient Bayesian Uncertainty Estimation for nnU-Net. In L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, & S. Li (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings* (pp. 535-544). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 13438 LNCS). Springer. https://doi.org/10.1007/978-3-031-16452-1_51

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Efficient Bayesian Uncertainty Estimation for nnU-Net

Zhao, Yidong; Yang, Changchun; Schweidtmann, Artur; Tao, Qian

DOI

[10.1007/978-3-031-16452-1_51](https://doi.org/10.1007/978-3-031-16452-1_51)

Publication date

2022

Document Version

Final published version

Published in

Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings

Citation (APA)

Zhao, Y., Yang, C., Schweidtmann, A., & Tao, Q. (2022). Efficient Bayesian Uncertainty Estimation for nnU-Net. In L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, & S. Li (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022 - 25th International Conference, Proceedings* (pp. 535-544). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 13438 LNCS). Springer. https://doi.org/10.1007/978-3-031-16452-1_51

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Efficient Bayesian Uncertainty Estimation for nnU-Net

Yidong Zhao¹, Changchun Yang¹, Artur Schweidtmann², and Qian Tao¹(✉)

¹ Department of Imaging Physics, Delft University of Technology,
Delft, The Netherlands
q.tao@tudelft.nl

² Department of Chemical Engineering, Delft University of Technology,
Delft, The Netherlands

Abstract. The self-configuring nnU-Net has achieved leading performance in a large range of medical image segmentation challenges. It is widely considered as the model of choice and a strong baseline for medical image segmentation. However, despite its extraordinary performance, nnU-Net does not supply a measure of uncertainty to indicate its possible failure. This can be problematic for large-scale image segmentation applications, where data are heterogeneous and nnU-Net may fail without notice. In this work, we introduce a novel method to estimate nnU-Net uncertainty for medical image segmentation. We propose a highly effective scheme for posterior sampling of weight space for Bayesian uncertainty estimation. Different from previous baseline methods such as Monte Carlo Dropout and mean-field Bayesian Neural Networks, our proposed method does not require a variational architecture and keeps the original nnU-Net architecture intact, thereby preserving its excellent performance and ease of use. Additionally, we boost the segmentation performance over the original nnU-Net via marginalizing multi-modal posterior models. We applied our method on the public ACDC and M&M datasets of cardiac MRI and demonstrated improved uncertainty estimation over a range of baseline methods. The proposed method further strengthens nnU-Net for medical image segmentation in terms of both segmentation accuracy and quality control.

Keywords: nnU-Net · Uncertainty estimation · Variational inference · Stochastic gradient descent

1 Introduction

Manually delineating the region of interest from medical images is immensely expensive in clinical applications. Among various automatic medical image segmentation methods, the U-Net architecture [24], in particular the self-configuring nnU-Net framework [14], has achieved state-of-the-art performance in a wide range of medical image segmentation tasks [2, 5]. Nevertheless, the reliability of neural networks remains a major concern for medical applications, due to their




Fig. 1. We observe that the network checkpoints at various training epochs make diverse predictions when the network is uncertain on a *hard* input $\mathbf{x}^h$ (a). On the contrary, predictions of good quality on an *easy* input $\mathbf{x}^e$ enjoys consistency across checkpoints (b). We leverage this phenomenon to perform Bayesian inference and quantify uncertainty of network predictions.

over-parametrization and poor interpretability. Using the softmax output as a categorical probability distribution proxy is known to cause model miscalibration, which often leads to over-confidence in errors [21]. This can be problematic for large-scale medical image segmentation applications, where data are heterogeneous and nnU-Net may fail without notice. Accurate uncertainty estimation is highly important in clinical deployment of automatic segmentation models [16].

Previous efforts on uncertainty estimation of neural networks are categorized as either Bayesian or non-Bayesian [21]. Several non-Bayesian strategies have been proposed to quantify the predictive uncertainty in medical image segmentation. Guo *et al.* proposed to use temperature scaling [12] to mitigate the miscalibration in softmax. Probabilistic U-Net [17] aims at generating a set of feasible segmentation maps. Baumgartner *et al.* proposed a hierarchical model to learn the distribution in latent spaces and make probabilistic predictions [1]. However, these methods require modification of the original network, sometimes to a significant extent, therefore hard to be integrated with the well-configured nnU-Net.

Unlike non-Bayesian approaches, Bayesian neural networks (BNNs) learn the weight posterior distribution given the training data. At inference time, a probabilistic prediction is made, which delivers a natural epistemic uncertainty [4]. However, deriving weight posterior is an intractable problem, and Variational Inference (VI) [3] is often used for its approximation. The popular mean-field VI [4] doubles the network parameter assuming independence among weights, and can be unstable during training [23]. Gal *et al.* proposed Monte-Carlo Dropout (MC-Dropout) [9] as a VI proxy by enabling dropout layers at test time. The method is theoretically and practically sound, however, recent works [6, 10, 11] reported silent failure and poor calibration, with degraded segmentation performance. Deep Ensembles [18] estimate network uncertainty via averaging independently trained networks. It is robust [23] but highly costly in computation, especially for models which demand lengthy training themselves.

We propose a novel method for medical image segmentation uncertainty estimation, which substantially outperforms the popular MC-Dropout, while being

significantly more efficient than deep ensembles. We are inspired by the optimization theory that during stochastic gradient descent (SGD), the network weights continuously explore the solution space, which is approximately equivalent to weight space posterior sampling [19, 20, 22]. While taken at appropriate moments of SGD, the network can be sampled *a posteriori*, and uncertainty can be reflected in the agreement among these posterior models, as illustrated in Fig. 1. This generic methodology enables VI approximation on any network, including the delicately-configured nnU-Net. The following contributions are made:

- A novel VI approximation method that realizes efficient posterior estimation of a deep model.
- An extension to the nnU-Net for uncertainty estimation in medical image segmentation, with significantly improved performance over several baseline methods including MC-Dropout.
- Further improvement of the segmentation performance beyond the original nnU-Net, as showcased on cardiac magnetic resonance (CMR) data.

2 Methods

2.1 Bayesian Inference

Given a training dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ with N image and segmentation label pairs, the BNN fits the posterior $p(\mathbf{w}|\mathcal{D})$ over network weights $\mathbf{w}$. At inference time, the probabilistic prediction on a test image $\mathbf{x}^*$ can be formulated as the marginalization over parameter $\mathbf{w}$:

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) = \int p(\mathbf{y}^*|\mathbf{x}^*, \mathbf{w})p(\mathbf{w}|\mathcal{D}) d\mathbf{w} \quad (1)$$

We describe the proposed trajectory-based posterior sampling from $p(\mathbf{w}|\mathcal{D})$ in Sect. 2.2 and our final multi-modal Bayesian inference based on cyclical training in Sect. 2.3.

2.2 SGD Bayesian Inference

Single-modal Posterior Sampling. During the SGD training, the weights first move towards a local optimum and then oscillate in the vicinity. The standard nnU-Net uses the polynomial learning rate decay scheme, in which the learning rate approaches zero as the epoch increases. The vanishing learning rate causes the network weights to converge. In order to make the weights explore a wider neighborhood around the local optimum, we follow [15] to set the learning rate to a constant value after γ fraction of the total epoch budget T . For each training epoch $t \in [\gamma T, T]$, the weight checkpoint $\mathbf{w}_t$ is approximately a sample drawn from a single-modal weight posterior: $\mathbf{w}_t \sim p(\mathbf{w}|\mathcal{D})$. In Fig. 1 (a) and (b), we show the exemplar predictions made by various checkpoints on an *easy* and a *hard* example, respectively.

Inference and Uncertainty Estimation. The $p(\mathbf{y}^*|\mathbf{x}^*, \mathbf{w})$ term in Eq. (1) corresponds to the model forward pass $f_{\mathbf{w}}(\mathbf{x})$ at weights $\mathbf{w}$. However, both the weight posterior and the integral are intractable for deep neural networks. In this work, we keep the intermediate weights (i.e. checkpoints) during SGD training as posterior samples after the learning stabilizes: $\mathbf{W} = \{\mathbf{w}_t | \gamma T < t \leq T\}$. Then, the Monte-Carlo approximation for the predictive posterior $p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})$ can be computed by taking the discrete version of Eq. (1):

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) \approx \frac{1}{n} \sum_{i=1}^n p(\mathbf{y}^*|\mathbf{x}^*, \mathbf{w}_{t_i}) \quad (2)$$

where $\mathbf{w}_{t_i} \in \mathbf{W}$ and n is the number of checkpoints out of the posterior sampling of the model. To generate the uncertainty map, we compute the entropy of the predictive C -class categorical distribution $H(\mathbf{y}_{i,j}^*)$ for each voxel (i, j) , with 0 indicating the lowest uncertainty.

$$H(\mathbf{y}_{i,j}^*) = - \sum_{k=1}^C p(\mathbf{y}_{i,j}^* = k | \mathbf{x}^*, \mathcal{D}) \log_2 p(\mathbf{y}_{i,j}^* = k | \mathbf{x}^*, \mathcal{D}) \quad (3)$$

Stochastic Weight Averaging. An alternative to the prediction space averaging in Eq. (2) is the Stochastic Weight Averaging (SWA) [15], which takes average in the weight space instead: $p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) = p(\mathbf{y}^*|\mathbf{x}^*, \bar{\mathbf{w}})$, where $\bar{\mathbf{w}} = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_{t_i}$, to construct a new aggregated model with improved performance [15]. After weight averaging, the network needs to perform one more epoch of forward passes to update the batch normalization layer parameters before inference.

2.3 Multi-modal Posterior Sampling

Cyclical Learning Rate. The traditional SGD converges to the neighborhood of a single local optimum, hence the checkpoints capture only a single mode of the weight posterior. The diversity in weight samples has proven beneficial for uncertainty estimation [7, 8]. In order to capture a multi-modal geometry in weight space, we employ a cyclical learning rate [13] which periodically drives the weights out of the attraction region of a local optimum. However, the cosine annealing scheme proposed in [13] causes instability in the nnU-Net training. We divide the training epoch budget T to M cycles, each of which consumes $T_c = \frac{T}{M}$ epochs. A high restart learning rate α_r is set only for the first epoch in each training cycle. After γ fraction of T_c epochs, we keep the gradient update step constant. Our proposed learning rate scheme is formulated as Eq. (4):

$$\alpha(t) = \begin{cases} \alpha_r, & t_c = 0 \\ \alpha_0 \left[1 - \frac{\min(t_c, \gamma T_c)}{T} \right]^\epsilon, & t_c > 0 \end{cases} \quad (4)$$

where $t_c = t \bmod T_c$ is the in-cycle epoch number and the constant exponent ϵ controls the decaying rate of the learning rate.

Multi-modal Checkpoint Ensemble. In the c^{th} training cycle, we collect checkpoints $\mathbf{W}_c = \{\mathbf{w}_t | \gamma T_c \leq t \leq T_c - 1\}$. Then the aggregated checkpoints of all M training cycles $\mathbf{W} = \cup_{c=1}^M \mathbf{W}_c$ consist of multi-modal posterior samples in the weight space. We build an ensemble of n models via sampling $\frac{n}{M}$ intermediate weights from each mode $\mathbf{W}_c$ for $c \in \{1, 2, \dots, M\}$. The multi-modal Bayesian inference is then performed as in Eq. (2). In comparison to [13], which collected one snapshot in each cycle, our method combines both local and global uncertainty of the weights.

3 Experiments

3.1 Experimental Setup

Datasets. We trained our network on the publicly available ACDC dataset¹ [2], which contains annotated end-diastolic (ED) and end-systolic (ES) volumes of 100 subjects. We randomly split 80% of The ACDC data for network training and kept the rest 20% to validate the in-domain (ID) performance. To evaluate the performance under a domain shift, i.e. out-of-domain (OOD), we also tested our method on the M&M dataset² [5] collected from different MRI vendors and medical centers. The annotated part includes 75 subjects collected from Vendor A (Siemens) in a single medical center and 75 scans by Vendor B (Phillips) from two medical centers.

Implementation Details. For CMR segmentation, we trained the 2D nnU-Net, with the proposed learning rate modulation in Sect. 2.2 and 2.3. During each training cycle, SGD with momentum was used for optimization with a batch size of 20. We set the initial learning rate as $\alpha_0 = 0.01$, the decay exponent $\epsilon = 0.9$ and the restart learning as $\alpha_r = 0.1$. We trained $M = 3$ cycles within a total training budget of $T = 1,200$ epochs, which is slightly higher than the standard nnU-Net which consumes 1,000 epochs for a single model training. The learning rate kept at a constant level after $\gamma = 80\%$ of the epochs in each individual training cycle. We chose the latest $n = 30$ checkpoints (epoch Step 2) in the last training cycle to build a single-modal checkpoints ensemble (*Single-modal Ckpt. Ens.*). For a multi-modal weight sampling (*Multi-modal Ckpt. Ens.*), we aggregated the latest 10 checkpoints of all three training cycles. We took the single model prediction at the end of the training as the baseline (*Vanilla*). Temperature scaling (*Temp. Scaling*) with $\tau = 1.5$ was implemented as a calibration baseline as in [12]. We also compared our proposed method to MC-Dropout [9] and deep ensemble [18] (*Deep Ens.*). The dropout probability was set as $p = 0.1$ and $n = 30$ predictions were drawn for uncertainty estimation. A 5-model deep ensemble was trained on the ACDC training data with random initialization. We also evaluated the performance of SWA [15], which averages the posterior weights sampled during the last training cycle, and we used the softmax predictions for uncertainty estimation.

¹ <https://www.creatis.insa-lyon.fr/Challenge/acdc>.

² <https://www.ub.edu/mnms>.




Fig. 2. (a) t-SNE plot of the weight space during SGD training. Dotted lines illustrate the transition between weight modes. (b) t-SNE plot of the posterior weights, which bounce in the neighborhood of different modes.

Evaluation Metrics. On each dataset, the segmentation performance was evaluated using the mean Dice coefficients of the three foreground classes, namely, left ventricle (LV), right ventricle (RV), and myocardium (MYO), over all volumes. We quantify the calibration performance through the expected calibration error (ECE) of voxels in the bounding box of all foreground classes in each volume [12, 21]. We divided the voxels into 100 bins according to their confidence from 0% to 100%. The ECE is defined as the weighted average of the confidence-accuracy difference in each bin: $ECE = \sum_{s=1}^{100} \frac{|B_s|}{N_v} |\text{conf}(B_s) - \text{acc}(B_s)|$, where B_s defines the set of voxels whose confidence falls in the range $((s-1)\%, s\%]$, $\text{conf}(B_s)$ is the mean confidence of voxels in bin B_s , $\text{acc}(B_s)$ represents the classification accuracy of the voxels in B_s and N_v is the total number of voxels.

3.2 Results

Training Trajectory Visualization. We first present the training trajectory in the weight space to validate our concept of multi-modal posterior. Figure 2(a) shows the t-SNE plot of the weights of the last decoder layer in nnU-Net during the three training cycles. We observe that the cyclical training triggered a long-range movement in the weight space such that multiple local optima were visited. In Fig. 2(b) we demonstrate both local and global weight posterior distribution. In each training cycle, the weights did not converge to a single point but kept bouncing in a local neighborhood, conforming to [19, 20]. The diversity in posterior modes is clearly shown by the three clusters in the t-SNE weight space. We note that similar behavior of multi-modality was also observed for posterior weights in other layers.

Segmentation and Calibration Performance. We list the average Dice coefficient of all methods in comparison in Table 1. The table demonstrates that except for MC-Dropout, which downgraded the performance slightly, the rest

Table 1. Dice coefficients on ID and OOD test sets.

Method	ACDC validation (ID)			M&M vendor A (OOD)			M&M vendor B (OOD)		
	RV	MYO	LV	RV	MYO	LV	RV	MYO	LV
Vanilla	0.911 $\pm$ 0.052	0.904 $\pm$ 0.026	0.944 $\pm$ 0.035	0.830 $\pm$ 0.119	0.810 $\pm$ 0.043	0.897 $\pm$ 0.063	0.878 $\pm$ 0.077	0.844 $\pm$ 0.054	0.894 $\pm$ 0.076
SWA	0.913 $\pm$ 0.054	0.910 $\pm$ 0.023	0.948 $\pm$ 0.033	0.856 $\pm$ 0.093	0.808 $\pm$ 0.041	0.896 $\pm$ 0.061	0.879 $\pm$ 0.067	0.845 $\pm$ 0.050	0.897 $\pm$ 0.069
MC-Dropout	0.906 $\pm$ 0.061	0.901 $\pm$ 0.028	0.940 $\pm$ 0.043	0.810 $\pm$ 0.146	0.810 $\pm$ 0.049	0.896 $\pm$ 0.066	0.880 $\pm$ 0.079	0.844 $\pm$ 0.058	0.891 $\pm$ 0.081
Deep Ens.	0.915 $\pm$ 0.051	0.912 $\pm$ 0.023	0.951 $\pm$ 0.029	0.857 $\pm$ 0.089	0.816 $\pm$ 0.041	0.902 $\pm$ 0.062	0.885 $\pm$ 0.068	0.849 $\pm$ 0.047	0.897 $\pm$ 0.064
Ckpt. Ens. (Single)	0.919 $\pm$ 0.051	0.912 $\pm$ 0.022	0.951 $\pm$ 0.032	0.851 $\pm$ 0.100	0.816 $\pm$ 0.041	0.902 $\pm$ 0.061	0.883 $\pm$ 0.070	0.850 $\pm$ 0.048	0.901 $\pm$ 0.065
Ckpt. Ens. (Multi)	0.918 $\pm$ 0.051	0.913 $\pm$ 0.024	0.951 $\pm$ 0.031	0.852 $\pm$ 0.105	0.818 $\pm$ 0.043	0.905 $\pm$ 0.060	0.885 $\pm$ 0.069	0.851 $\pm$ 0.050	0.899 $\pm$ 0.071

methods improved the segmentation performance in comparison to the Vanilla nnU-Net. The proposed checkpoint ensembles achieved overall the best performance except for the RV on Vendor A.

Table 2. ECE (%) on ID and OOD test sets.

Methods	ACDC validation (ID)	M&M vendor A (OOD)	M&M vendor B (OOD)
Vanilla	2.56	4.34	3.79
Temp. scaling	2.18	3.91	3.46
SWA	2.39	4.07	3.70
MC-dropout	1.70	3.41	2.95
Deep Ens	1.63	3.16	2.85
Ckpt. Ens. (Single)	1.25	2.83	2.69
Ckpt. Ens. (Multi)	1.25	2.75	2.61

The ECE scores are listed in Table 2. From the table, we can observe that the temperature scaling slightly improved the calibration performance. The results of Vanilla and SWA show that a single model was poorly calibrated compared to all 4 Bayesian methods. MC-Dropout improved calibration but was not as effective as Deep Ensemble. The proposed single-model checkpoint ensemble achieved the best performance on ID data, and the multi-modal checkpoint ensemble further improved the performance on OOD data.

Qualitative Results. Some representative examples of the segmentation and uncertainty maps are shown in Fig. 3. For the successful segmentation in Fig. 3(a), all methods highlighted the prediction borders as uncertain compared to regions, which conforms to our intuition in practice that different observers may have slightly different border delineations. Figure 3(b–d) show three cases in which the network failed, mainly in the RV region. In practice, this is the most common failure even with the state-of-the-art nnU-Net. We can observe that MC-Dropout frequently output low uncertainty in the area of wrong predictions (RV). Deep Ensemble correctly identified the uncertain areas in case (b) and (d). However, the limited number of models in Deep Ensemble may miss part of




Fig. 3. Predictions (Pred.) and estimated uncertainty maps (Uncert.) on a successful case (a) and three partially failed cases (b–d). In case (a) all the methods highlighted the border as uncertain. MC-Dropout failed in all the three cases (b–d) to delineate RV, while reporting low uncertainty in the corresponding area. Deep Ensemble is robust but missed part of the uncertain areas in case (c). Multi-modal weight sampling detected the failed RV area more robustly than the single-modal version in case (d). (Color figure online)

the error, as shown in Fig. 3(c) (red arrow). We further demonstrate the benefit of multi-modality in weight space through Fig. 3(d) which shows an apical slice with a small RV. In this case, only deep ensemble and the multi-modal checkpoints ensemble successfully detected the uncertainty in RV region, while other methods just output high confidence, indicating “silent” failure, i.e. mistake that will escape notice in automatic segmentation pipelines.

In all three challenging cases, our proposed method robustly detected the nnU-Net failure on RV. Interestingly, we note that the estimated uncertainty map strongly correlated with true RV area, which is highly beneficial and informative for further manual contour correction in clinical practice.

4 Conclusion

In this work, we proposed an efficient Bayesian inference approximation method for nnU-Net, from the SGD training perspective. Our method does not require adapting the nnU-Net configuration and is highly efficient in computation. Our experimental results on both in-domain and out-of-domain cardiac MRI data proved its effectiveness, in comparison with established baseline methods such as MC-Dropout and Deep Ensemble. The proposed uncertainty estimation further strengthens nnU-Net for practical applications of medical image segmentation, in terms of both *uncertainty estimation* and *segmentation performance*.

Acknowledgement. The authors gratefully acknowledge TU Delft AI Initiative for financial support.

References

1. Baumgartner, C.F., Tezcan, K.C., Chaitanya, K., Hötter, A.M., Muehlematter, U.J., Schawkat, K., Becker, A.S., Donati, O., Konukoglu, E.: PHiSeg: capturing uncertainty in medical image segmentation. In: Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.-T., Khan, A. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 119–127. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_14
2. Bernard, O., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE Trans. Med. Imaging **37**(11), 2514–2525 (2018)
3. Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: a review for statisticians. J. Am. Stat. Assoc. **112**(518), 859–877 (2017)
4. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: International Conference on Machine Learning, pp. 1613–1622. PMLR (2015)
5. Campello, V.M., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: the M&MS challenge. IEEE Trans. Med. Imaging **40**(12), 3543–3554 (2021)
6. Folgoc, L.L., et al.: Is MC dropout bayesian? arXiv preprint [arXiv:2110.04286](https://arxiv.org/abs/2110.04286) (2021)
7. Fort, S., Hu, H., Lakshminarayanan, B.: Deep ensembles: a loss landscape perspective. arXiv preprint [arXiv:1912.02757](https://arxiv.org/abs/1912.02757) (2019)
8. Fuchs, M., Gonzalez, C., Mukhopadhyay, A.: Practical uncertainty quantification for brain tumor segmentation (2021)
9. Gal, Y., Ghahramani, Z.: Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: International Conference on Machine Learning, pp. 1050–1059. PMLR (2016)
10. Gonzalez, C., Gotkowsky, K., Bucher, A., Fischbach, R., Kaltenborn, I., Mukhopadhyay, A.: Detecting when pre-trained nnU-Net models fail silently for Covid-19 lung lesion segmentation. In: de Bruijne, M., et al. (eds.) MICCAI 2021. LNCS, vol. 12907, pp. 304–314. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87234-2_29
11. Gonzalez, C., Mukhopadhyay, A.: Self-supervised out-of-distribution detection for cardiac CMR segmentation. In: Medical Imaging with Deep Learning (2021)

12. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: International Conference on Machine Learning, pp. 1321–1330. PMLR (2017)
13. Huang, G., Li, Y., Pleiss, G., Liu, Z., Hopcroft, J.E., Weinberger, K.Q.: Snapshot ensembles: train 1, get M for free. arXiv preprint [arXiv:1704.00109](https://arxiv.org/abs/1704.00109) (2017)
14. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**(2), 203–211 (2021)
15. Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., Wilson, A.G.: Averaging weights leads to wider optima and better generalization. arXiv preprint [arXiv:1803.05407](https://arxiv.org/abs/1803.05407) (2018)
16. Jungo, A., Reyes, M.: Assessing reliability and challenges of uncertainty estimations for medical image segmentation. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11765, pp. 48–56. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32245-8_6
17. Kohl, S.A., et al.: A probabilistic U-Net for segmentation of ambiguous images. arXiv preprint [arXiv:1806.05034](https://arxiv.org/abs/1806.05034) (2018)
18. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. arXiv preprint [arXiv:1612.01474](https://arxiv.org/abs/1612.01474) (2016)
19. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for Bayesian uncertainty in deep learning. In: Advances in Neural Information Processing Systems 32 (2019)
20. Mandt, S., Hoffman, M.D., Blei, D.M.: Stochastic gradient descent as approximate Bayesian inference. arXiv preprint [arXiv:1704.04289](https://arxiv.org/abs/1704.04289) (2017)
21. Mehrtash, A., Wells, W.M., Tempany, C.M., Abolmaesumi, P., Kapur, T.: Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Trans. Med. Imaging* **39**(12), 3868–3878 (2020)
22. Mingard, C., Valle-Pérez, G., Skalse, J., Louis, A.A.: Is SGD a Bayesian sampler? well, almost. *J. Mach. Learn. Res.* **22** (2021)
23. Ovadia, Y., et al.: Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. arXiv preprint [arXiv:1906.02530](https://arxiv.org/abs/1906.02530) (2019)
24. 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