



Knowing What You Know in Brain Segmentation Using Bayesian Deep Neural Networks

Patrick McClure^{1,2*}, Nao Rho^{1,2}, John A. Lee^{2,3}, Jakub R. Kaczmarzyk⁴, Charles Y. Zheng^{1,2}, Satrajit S. Ghosh⁴, Dylan M. Nielson^{2,3}, Adam G. Thomas³, Peter Bandettini² and Francisco Pereira^{1,2}

¹ Machine Learning Team, National Institute of Mental Health, Bethesda, MD, United States, ² Section on Functional Imaging Methods, National Institute of Mental Health, Bethesda, MD, United States, ³ Data Sharing and Science Team, National Institute of Mental Health, Bethesda, MD, United States, ⁴ McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA, United States

In this paper, we describe a Bayesian deep neural network (DNN) for predicting FreeSurfer segmentations of structural MRI volumes, in minutes rather than hours. The network was trained and evaluated on a large dataset ($n = 11,480$), obtained by combining data from more than a hundred different sites, and also evaluated on another completely held-out dataset ($n = 418$). The network was trained using a novel spike-and-slab dropout-based variational inference approach. We show that, on these datasets, the proposed Bayesian DNN outperforms previously proposed methods, in terms of the similarity between the segmentation predictions and the FreeSurfer labels, and the usefulness of the estimate uncertainty of these predictions. In particular, we demonstrated that the prediction uncertainty of this network at each voxel is a good indicator of whether the network has made an error and that the uncertainty across the whole brain can predict the manual quality control ratings of a scan. The proposed Bayesian DNN method should be applicable to any new network architecture for addressing the segmentation problem.

Keywords: brain segmentation, deep learning, magnetic resonance imaging, Bayesian neural networks, variational inference, automated quality control

OPEN ACCESS

Edited by:

Roberto Toro,
Institut Pasteur, France

Reviewed by:

Danilo Bzdok,
Julich Research Centre, Germany
Nicola Amoroso,
University of Bari Aldo Moro, Italy

*Correspondence:

Patrick McClure
patrick.mcclure@nih.gov

Received: 20 June 2019

Accepted: 18 September 2019

Published: 17 October 2019

Citation:

McClure P, Rho N, Lee JA, Kaczmarzyk JR, Zheng CY, Ghosh SS, Nielson DM, Thomas AG, Bandettini P and Pereira F (2019) Knowing What You Know in Brain Segmentation Using Bayesian Deep Neural Networks. *Front. Neuroinform.* 13:67. doi: 10.3389/fninf.2019.00067

1. INTRODUCTION

Identifying which voxels in a structural magnetic resonance imaging (sMRI) volume correspond to different brain structures (i.e., segmentation) is an essential processing step in neuroimaging analyses. These segmentations are often generated using the FreeSurfer package (Fischl, 2012), a process which can take a day or more for each subject (FreeSurfer, 2018). The computational resources for doing this at a scale of hundreds to thousands of subjects are beyond the capabilities of the computational resources available to most researchers. This has led to an interest in the use of deep neural networks as a general approach for learning to predict the outcome of a processing task, given the input data, in a much shorter time period than the processing would normally take. In particular, several deep neural networks have been trained to perform segmentation of brain sMRI volumes (Ronneberger et al., 2015; Fedorov et al., 2017a,b; Li et al., 2017; Dolz et al., 2018; Roy et al., 2018b), taking between a few seconds and a few minutes per volume. These networks predict a manual or an automated segmentation from the structural volumes [Fedorov et al. (2017a,b), Dolz et al. (2018), and Roy et al. (2018b) used FreeSurfer, Petersen et al. (2010) used GIF (Cardoso et al., 2015), and Rajchl et al. (2018) used FSL (Jenkinson et al., 2012), SPM (Friston et al., 1994) and MALP-EM (Ledig et al., 2015)].

These networks, however, have been trained on a limited number (on the order of hundreds or low thousands) of examples from a limited number of sites (i.e., locations and/or scanners), which can lead to poor cross-site generalization for complex segmentation tasks with a large number of classes (McClure et al., 2018). This includes several of the recent DNNs proposed for fine-grain sMRI segmentation. (Note: We focus on DNNs which predict >30 classes.)

Roy et al. (2018b) performed 33 class segmentation using 581 sMRI volumes from the IXI dataset to train an initial model and then fine-tuned on 28 volumes from the MALC dataset (Marcus et al., 2007). They showed an ~9.4% average Dice loss on out-of-site data from the ADNI-29 (Mueller et al., 2005), CANDI (Kennedy et al., 2012), and IBSR (Rohlfing, 2012) datasets. Fedorov et al. (2017a) used 770 sMRI volumes from HCP (Van Essen et al., 2013) to train an initial model and then fine-tuned on 7 volumes from the FBIRN dataset (Keator et al., 2016). Li et al. (2017) performed a 160 class segmentation using 443 sMRI volumes from the ADNI dataset (Peterson et al., 2010) for training. Rajchl et al. (2018) used 5,000 sMRI volumes from the UK Biobank Sudlow et al. (2015) dataset to train a model by using multi-task learning to simultaneously predict 4 SPM labels, 17 FSL labels, and 139 MALP-EM labels. However, Fedorov et al. (2017a), Li et al. (2017), and Rajchl et al. (2018) did not report test results for sites that were not used during training.

These results show that it is possible to train a neural network to carry out segmentation of a sMRI volume. However, they provide a limited indication of whether such a network would work on data from any new site not encountered in training. While fine-tuning on labeled data from new sites can improve performance, even while using small amounts of data (Fedorov et al., 2017a; McClure et al., 2018; Roy et al., 2018b), a robust neural network segmentation tool should generalize to new sites without any further effort.

As part of the process of adding segmentation capabilities to the “Nobrainier” tool¹, we trained a network to predict FreeSurfer segmentations given a training set of ~10,000 sMRI volumes. This paper describes this process, as well as a quantitative and qualitative evaluation of the performance of the resulting model.

Beyond the segmentation performance of the network, a second aspect of interest to us is to understand whether it is feasible for a network to indicate how confident it is about its prediction at each location in the brain. We expect the network to make errors, be it because of noise, unusual positioning of the brain, very different contrast than what it was trained on, etc. Because our model is probabilistic and seeks to learn uncertainties, we expect it to be less confident in its predictions in such cases. It is also possible that, for certain locations, there are very similar brain structures labeled as different regions in different people. In such locations, there would be a limit to how well the network could perform, the Bayes error rate (Hastie et al., 2005). Additionally, the network should be less confident for examples that are very different from those seen in the training set (e.g., contain large artifacts). While prediction uncertainty can be computed for standard neural networks,

as done by Dolz et al. (2018), these uncertainty estimates are often overconfident (Guo et al., 2017; McClure and Kriegeskorte, 2017). Bayesian neural networks (BNNs) have been proposed as a solution to this issue. One popular BNN approach is Monte-Carlo (MC) Bernoulli Dropout (Srivastava et al., 2014; Gal and Ghahramani, 2016). Using this method, Li et al. (2017) and Roy et al. (2018a) showed that the segmentation performance of the BNN predictions was better for voxels with low dropout sampling-based uncertainties and that injected input noise can lead to increased uncertainty. Roy et al. (2018a) also found that using MC Bernoulli dropout decreased the drop in segmentation performance from 9.4% to 7.8% on average when testing on data from new sites compared to Roy et al. (2018b). However, MC Bernoulli dropout does not learn dropout probabilities from data, which can lead to not properly modeling the uncertainty of the predicted segmentation. Recent works have shown that these dropout probabilities can be learned using a concrete relaxation (Gal et al., 2017). Additionally, learning individual uncertainties for each weight has been shown to be beneficial for many purposes (e.g., pruning and continual learning) (Blundell et al., 2015; McClure et al., 2018; Nguyen et al., 2018). In this paper, we propose using both learned dropout uncertainties and individual weight uncertainties.

Finally, we test the hypothesis that overall prediction uncertainty across an entire image reflects its “quality,” as measured by human quality control (QC) scores. Given the effort required to produce such scores, there have been multiple attempts to either crowdsource the process (Keshavan et al., 2018) or automate it (Esteban et al., 2017). The latter, in particular, does not rely on segmentation information, so we believe it is worthwhile to test whether uncertainty derived from segmentation is more effective.

2. METHODS

2.1. Data

2.1.1. Imaging Datasets

We combined several datasets (Table 1), many which themselves contain data from multiple sites, into a single dataset with 11,480 T1 sMRI volumes. In-site validation and test sets were created from the combined dataset using a 80-10-10 training-validation-test split. This resulted in a training set of 9,184 volumes, a validation set of 1,148 volumes, and a test set of 1,148 volumes. The training set was used for training the networks, the validation set for setting DNN hyperparameters (e.g., Bernoulli dropout probabilities), and the test set was used for evaluating the performance of the DNNs on new data from the same sites that were used for training.

We additionally used 418 sMRI volumes from the NNDSP dataset (Lee et al., 2018) as a held-out dataset to test generalization of the network to an unseen site. In addition to sMRI volumes, each NNDSP sMRI volume was given a QC score from 1 to 4, higher scores corresponding to worse scan quality, by two raters (3 if values differed by more than 1), as described in Blumenthal et al. (2002). If a volume had a QC score greater than 2, it was labeled as a bad quality scan; otherwise, the scan was labeled as a good quality scan.

¹<https://github.com/neuronets/nobrainier>

TABLE 1 | The number of examples used from different datasets.

Dataset	Number of examples
CoRR (Zuo et al., 2014)	3,039
OpenfMRI (Poldrack et al., 2013)	1,873
NKI (Nooner et al., 2012)	1,136
SLIM (Liu et al., 2017)	1,003
ABIDE (Di Martino et al., 2014)	992
HCP (Van Essen et al., 2013)	956
ADHD200 (Bellec et al., 2017)	719
CMI (Alexander et al., 2017)	611
SALD (Wei et al., 2018)	477
Buckner (Biswal et al., 2010)	183
HBNSI (O'Connor et al., 2017)	178
GSP (Holmes et al., 2015)	152
Haxby (Haxby et al., 2011; Nastase et al., 2017)	55
Gobbini (di Oleggio Castello et al., 2017)	51
ICBM (Mazziotta et al., 2001)	45
Barrios (Vzquez et al., 2016)	10

2.1.2. Segmentation Target

We computed 50-class FreeSurfer (Fischl, 2012) segmentations, as in Fedorov et al. (2017a), for all subjects in each of the datasets described earlier. These were used as the labels for prediction. Although, FreeSurfer segmentations may not be perfectly correct, as compared to manual, expert segmentations, using them allowed us to create a large training dataset, as one could not feasibly label it by hand. FreeSurfer trained networks can also outperform FreeSurfer segmentations when compared to expert segmentations (Roy et al., 2018b). The trained network could be fine-tuned with expert small amounts of labeled data, which would likely improve the results (McClure et al., 2018; Roy et al., 2018b).

2.1.3. Data Pre-processing

The sMRI volumes were resampled to 1mm isotropic cubic volumes of 256 voxels per side and the voxel intensities were normalized according to Freesurfer's mri_convert with the conform flag. After resampling, input volumes were individually z-scored across voxels. We then split each sMRI volume into 512 non-overlapping $32 \times 32 \times 32$ sub-volumes, similarly to Fedorov et al. (2017b) and Fedorov et al. (2017a), to be used as inputs for the neural network. The prediction target is the corresponding segmentation sub-volume. This resulted in 512 pairs (\mathbf{x}, \mathbf{y}) , of sMRI and label sub-volumes, respectively, for each sMRI volume.

2.2. Convolutional Neural Network

2.2.1. Architecture

Several deep neural network architectures have been proposed for brain segmentation, such as U-net (Ronneberger et al., 2015), QuickNAT (Roy et al., 2018b), HighResNet (Li et al., 2017), and MeshNet (Fedorov et al., 2017a,b). We chose MeshNet because of its relatively simple structure, its lower number of learned parameters, and its competitive performance, since the computational cost of Bayesian neural

TABLE 2 | The MeshNet dilated convolutional neural network architecture used for brain segmentation.

Layer	Filter	Padding	Dilation (l)	Non-linearity
1	96×3^3	1	1	ReLU
2	96×3^3	1	1	ReLU
3	96×3^3	1	1	ReLU
4	96×3^3	2	2	ReLU
5	96×3^3	4	4	ReLU
6	96×3^3	8	8	ReLU
7	96×3^3	1	1	ReLU
8	50×1^3	0	1	Softmax

networks scales based on structural complexity and number of parameters.

MeshNet uses dilated convolutional layers (Yu and Koltun, 2015) due to the 3D structural nature of sMRI data. Applying a discrete volumetric dilated convolutional layer to one input channel for one weight filter can be expressed as:

$$\begin{aligned}
 (\mathbf{w}_f * \mathbf{h})_{i,j,k} &= \sum_{\tilde{i}=-a}^a \sum_{\tilde{j}=-b}^b \sum_{\tilde{k}=-c}^c w_{f,\tilde{i},\tilde{j},\tilde{k}} h_{i-\tilde{i},j-\tilde{j},k-\tilde{k}} = (\mathbf{w}_f * \mathbf{h})_{\mathbf{v}} \\
 &= \sum_{\mathbf{t} \in \mathcal{W}_{abc}} w_{f,\mathbf{t}} h_{\mathbf{v}-\mathbf{t}}.
 \end{aligned} \quad (1)$$

where h is the input to the layer, a , b , and c are the bounds for the i , j , and k axes of the filter with weights $w_f(i, j, k)$ is the voxel, \mathbf{v} , where the convolution is computed. The set of indices for the elements of \mathbf{w}_f can be defined as $\mathcal{W}_{abc} = \{-a, \dots, a\} \times \{-b, \dots, b\} \times \{-c, \dots, c\}$. The dilation factor, number of filters, and other details of the MeshNet-like architecture that we used for all experiments is shown in **Table 2**. Note that we increased the number of filters per layer from 72 to 96, compared to Fedorov et al. (2017a) and McClure et al. (2018), since we greatly increased the number of training volumes.

2.2.2. Maximum a Posteriori Estimation

When training a neural network, the weights of the network, \mathbf{w} , are often learned using maximum likelihood estimation (MLE). For MLE, $\log p(\mathcal{D}|\mathbf{w})$ is maximized where $\mathcal{D} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$ is the training dataset and $(\mathbf{x}_n, \mathbf{y}_n)$ is the n th input-output example. This often overfits, however, so we used a prior on the network weights, $p(\mathbf{w})$, to obtain a maximum a posteriori (MAP) estimate, by optimizing $\log p(\mathbf{w}|\mathcal{D})$:

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \sum_{n=1}^N \log p(\mathbf{y}_n | \mathbf{x}_n, \mathbf{w}) + \log p(\mathbf{w}). \quad (2)$$

We used a fully factorized Gaussian prior (i.e., $p(w_{f,\tilde{i},\tilde{j},\tilde{k}}) = \mathcal{N}(0, 1)$). This results in the MAP weights being learned by minimizing the softmax cross-entropy with L2 regularization. At test time, this point estimate approximation, \mathbf{w}^* , is used to make a prediction for new examples:

$$p(\mathbf{y}_{test} | \mathbf{x}_{test}) \approx p(\mathbf{y}_{test} | \mathbf{x}_{test}, \mathbf{w}^*) \quad (3)$$

2.2.3. Approximate Bayesian Inference

In Bayesian inference for neural networks, a distribution of possible weights is learned instead of just a MAP point estimate. Using Bayes' rule, $p(\mathbf{w}|\mathcal{D}) = p(\mathcal{D}|\mathbf{w})p(\mathbf{w})/p(\mathcal{D})$, where $p(\mathbf{w})$ is the prior over weights. However, directly computing the posterior, $p(\mathbf{w}|\mathcal{D})$, is often intractable, particularly for DNNs. As a result, an approximate inference method must be used.

One of the most popular approximate inference methods for neural networks is variational inference, since it scales well to large DNNs. In variational inference, the posterior distribution $p(\mathbf{w}|\mathcal{D})$ is approximated by a learned variational distribution of weights $q_\theta(\mathbf{w})$, with learnable parameters θ . This approximation is enforced by minimizing the Kullback-Leibler divergence (KL) between $q_\theta(\mathbf{w})$, and the true posterior, $p(\mathbf{w}|\mathcal{D})$, $\text{KL}[q_\theta(\mathbf{w})||p(\mathbf{w}|\mathcal{D})]$, which measures how $q_\theta(\mathbf{w})$ differs from $p(\mathbf{w}|\mathcal{D})$ using relative entropy. This is equivalent to maximizing the variational lower bound (Hinton and Van Camp, 1993; Graves, 2011; Blundell et al., 2015; Kingma et al., 2015; Gal and Ghahramani, 2016; Louizos and Welling, 2017; Molchanov et al., 2017), also known as the evidence lower bound (ELBO),

$$\mathcal{L}_{ELBO}(\theta) = \mathcal{L}_{\mathcal{D}}(\theta) - \mathcal{L}_{KL}(\theta), \tag{4}$$

where $\mathcal{L}_{\mathcal{D}}(\theta)$ is

$$\mathcal{L}_{\mathcal{D}}(\theta) = \sum_{n=1}^N \mathbb{E}_{q_\theta(\mathbf{w})}[\log p(\mathbf{y}_n|\mathbf{x}_n, \mathbf{w})] \tag{5}$$

and $\mathcal{L}_{KL}(\theta)$ is the KL divergence between the variational distribution of weights and the prior,

$$\mathcal{L}_{KL}(\theta) = \text{KL}[q_\theta(\mathbf{w})||p(\mathbf{w})], \tag{6}$$

which measures how $q_\theta(\mathbf{w})$ differs from $p(\mathbf{w})$ using relative entropy.

Maximizing $\mathcal{L}_{\mathcal{D}}$ seeks to learn a $q_\theta(\mathbf{w})$ that explains the training data, while minimizing \mathcal{L}_{KL} (i.e., keeping $q_\theta(\mathbf{w})$ close to $p(\mathbf{w})$) prevents learning a $q_\theta(\mathbf{w})$ that overfits to the training data.

The objective function in Equation (4) is usually impractical to compute for deep neural networks, due to both: (1) being a full-batch approach and (2) integrating over $q_\theta(\mathbf{w})$. (1) is often dealt with by using stochastic mini-batch optimization (Robbins and Monro, 1951) and (2) is often approximated using Monte Carlo sampling. As discussed in Graves (2011) and Kingma et al. (2015), these methods can be used to perform stochastic gradient variational Bayes (SGVB) in deep neural networks. For each parameter update, an unbiased estimate of $\nabla_\theta \mathcal{L}_{\mathcal{D}}$ for a mini-batch, $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_M, \mathbf{y}_M)\}$, is calculated using one weight sample, \mathbf{w}_m , from $q_\theta(\mathbf{w})$ for each mini-batch example. This results in the following approximation to Equation (4):

$$\mathcal{L}_{ELBO}(\theta) \approx \mathcal{L}_{\mathcal{D}}^{SGVB}(\theta) - \mathcal{L}_{KL}(\theta), \tag{7}$$

where

$$\mathcal{L}_{\mathcal{D}}(\theta) \approx \mathcal{L}_{\mathcal{D}}^{SGVB}(\theta) = \frac{N}{M} \sum_{m=1}^M \log p(\mathbf{y}_m|\mathbf{x}_m, \mathbf{w}_m). \tag{8}$$

At test time, the weights, \mathbf{w} would ideally be marginalized out, $p(\mathbf{y}_{test}|\mathbf{x}_{test}) = \int p(\mathbf{y}_{test}|\mathbf{x}_{test}, \mathbf{w})q_\theta(\mathbf{w})d\mathbf{w}$, when making a prediction for a new example. However, this is often impractical to compute for DNNs, so a Monte-Carlo approximation is often used. This results in the prediction of a new example being made by averaging the predictions of multiple weight samples from $q_\theta(\mathbf{w})$ (Figure 1):

$$p(\mathbf{y}_{test}|\mathbf{x}_{test}) \approx \frac{1}{N_{MC}} \sum_n^{N_{MC}} p(\mathbf{y}_{test}|\mathbf{x}_{test}, \mathbf{w}_n) \tag{9}$$

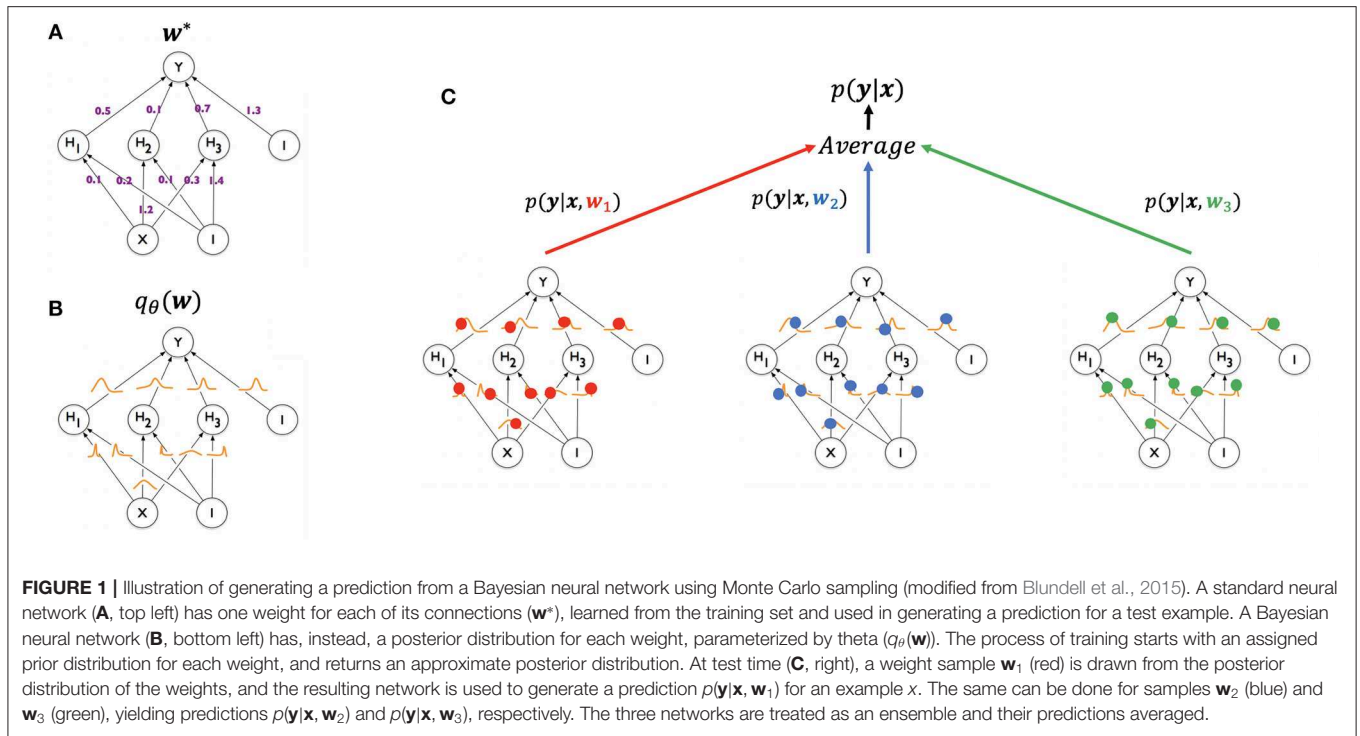
where $\mathbf{w}_n \sim q_\theta(\mathbf{w})$.

2.2.3.1. MC Bernoulli dropout

For MC Bernoulli dropout (BD) (Gal and Ghahramani, 2016), we drew weights from $q_\theta(\mathbf{w})$ by drawing a Bernoulli random variable ($b_{i,j,k} \sim \text{Bern}(p_l)$), where i, j, k are the indices of the volume axes, for every element of the layer, l , input, \mathbf{h} , and then elementwise multiplying \mathbf{b} and \mathbf{h} before applying the next dilated convolutional layer. This effectively sets the filter weights to zero when applied to a dropped element. Gal and Ghahramani (2016) approximated the KLD between this Bernoulli variational distribution and a zero-mean Gaussian by replacing the variational distribution with a mixture of Gaussians, resulting in an L2-like penalty. However, this can lead to pathological behavior due to Bernoulli distributions not having support over all real numbers (Hron et al., 2018). In Bernoulli dropout, p_l codes for the uncertainty of the weights and is often set layerwise via hyperparameter search (for our experiments, we found the best value of p to be 0.9 after searching over the values of 0.95, 0.9, 0.75, and 0.5 using the validation set). However, Bayesian models would ideally learn how uncertain to be for each weight.

2.2.3.2. Spike-and-slab dropout with learned model uncertainty

We propose a form of dropout that both learns the dropout probability for each filter using a concrete relaxation of dropout (Gal et al., 2017), and an individual uncertainty for each weight using fully factorized Gaussian (FFG) filters (Graves, 2011; Blundell et al., 2015; Molchanov et al., 2017; McClure et al., 2018; Nguyen et al., 2018). This is in contrast to previous spike-and-slab dropout methods, which did not learn the model (or epistemic) uncertainty (Der Kiureghian and Ditlevsen, 2009; Kendall and Gal, 2017) from data either by learning the dropout probabilities or by learning the variance parameter of the Gaussian components of the weights (McClure and Kriegeskorte, 2017). In our proposed method, we assume each of the F filters are independent (i.e., $p(\mathbf{w}) = \prod_{f=1}^F p(\mathbf{w}_f)$), as done in previous FFG methods (Graves, 2011; Blundell et al., 2015; Molchanov et al., 2017; McClure et al., 2018; Nguyen et al., 2018). We then decompose each filter into a dropout-based component, b_f , and a Gaussian component, \mathbf{g}_f , such that $\mathbf{w}_f = b_f \mathbf{g}_f$. Per this decomposition, we perform variational inference on the joint distribution of $\{b_1, \dots, b_F, \mathbf{g}_1, \dots, \mathbf{g}_F\}$, instead of on $p(\mathbf{w})$ directly (Titsias and Lázaro-Gredilla, 2011; McClure and Kriegeskorte, 2017). We then assume each element of \mathbf{g}_f is independent (i.e.,



$p(\mathbf{g}_f) = \prod_{t \in \mathcal{W}_{abc}} p(\mathbf{g}_{f,t})$, and that each weight is Gaussian (i.e., $\mathbf{g}_{f,t} \sim \mathcal{N}(\mu_{f,t}, \sigma_{f,t}^2)$) with learned parameters $\mu_{f,t}$ and $\sigma_{f,t}$. Instead of drawing each b_f from $Bern(p_f)$, we draw them from a concrete distribution (Gal et al., 2017) with a learned dropout probability, p_f , for each filter:

$$b_f = \text{sigmoid}\left(\frac{1}{t}(\log p_f - \log(1 - p_f) + \log u - \log(1 - u))\right) \quad (10)$$

where $u \sim \text{Unif}(0, 1)$. This concrete distribution converges to the Bernoulli distribution as the sigmoid scaling parameter, t , goes to zero. (In this paper, we used $t = 0.02$). As discussed in Kingma et al. (2015) and Molchanov et al. (2017), randomly sampling each $\mathbf{g}_{f,t}$ for each mini-batch example can be computationally expensive, so we used the fact that the sum of independent Gaussian variables is also Gaussian to move the noise from the weights to the convolution operation, as in McClure et al. (2018). For, dilated convolutions and the proposed spike-and-slab variational distribution, this is described by:

$$(\mathbf{w}_f * \mathbf{h})_v = b_f(\mathbf{g}_f * \mathbf{h})_v \quad (11)$$

where

$$(\mathbf{g}_f * \mathbf{h})_v \sim \mathcal{N}(\mu_{f,v}^*, (\sigma_{f,v}^*)^2), \quad (12)$$

$$\mu_{f,v}^* = \sum_{t \in \mathcal{W}_{abc}} \mu_{f,t} h_{v-t}, \quad (13)$$

and

$$(\sigma_{f,v}^*)^2 = \sum_{t \in \mathcal{W}_{abc}} \sigma_{f,t}^2 h_{v-t}^2. \quad (14)$$

For this spike-and-slab dropout (SSD) implementation, we used a spike-and-slab prior, instead of the Gaussian prior used by Gal and Ghahramani (2016) and Gal et al. (2017). Using a spike-and-slab prior with MC Bernoulli dropout was discussed in Gal (2016), but not implemented. As in the variational distribution, each filter is independent in the prior. Per the spike-and-slab decomposition discussed above, the KL-divergence term of the ELBO can be written as

$$\mathcal{L}_{KL}(\theta) = \sum_{f=1}^F \text{KL}[q_{p_f}(b_f)q_{\mu,\sigma}(\mathbf{g}_f) || p(b_f)p(\mathbf{g}_f)], \quad (15)$$

where $\theta = \bigcup_f \bigcup_{t \in \mathcal{W}_{abc}} \{p_f, \mu_{f,t}, \sigma_{f,t}\}$ are the learned parameters and $p(b_f)$ and $p(\mathbf{g}_f)$ are priors. Assuming that each weight in a filter is independent, as commonly done in the literature (Graves, 2011; Blundell et al., 2015; Nguyen et al., 2018), allows the term to be rewritten as

$$\mathcal{L}_{KL}(\theta) = \sum_{f=1}^F (\text{KL}[q_{p_f}(b_f) || p(b_f)] + \sum_{t \in \mathcal{W}_{abc}} \text{KL}[q_{\mu,\sigma}(\mathbf{g}_{f,t}) || p(\mathbf{g}_{f,t})]). \quad (16)$$

For $\text{KL}[q_{p_f} || p(b_f)]$, we used the KL-divergence between two Bernoulli distributions,

$$\text{KL}[q_{p_f}(b_f) || p(b_f)] = p_f \log \frac{p_f}{p_{\text{prior}}} + (1 - p_f) \log \frac{1 - p_f}{1 - p_{\text{prior}}}, \quad (17)$$

since we used a relatively small sigmoid scaling parameter. Using $p(g_{f,t}) = \mathcal{N}(\mu_{prior}, \sigma_{prior}^2)$,

$$KL[q_{\mu,\sigma}(g_{f,t})||p(g_{f,t})] = \log \frac{\sigma_{prior}}{\sigma_{f,t}} + \frac{\sigma_{f,t}^2 + (\mu_{f,t} - \mu_{prior})^2}{2\sigma_{prior}^2} - \frac{1}{2}. \tag{18}$$

For this paper, the spike-and-slab prior parameters were set as $p_{prior} = 0.5, \mu_{prior} = 0$, and $\sigma_{prior} = 0.1$. $p_{prior} = 0.5$ corresponds to a maximum entropy prior (i.e., in the absence of new data be maximally uncertain). Alternatively, a p_{prior} close to 0 is a sparsity prior (i.e., in the absence of data do not use a filter).

2.3. Implementation Details

The DNNs were implemented using Tensorflow (Abadi et al., 2016). During training, the parameters of each DNN were updated using Adam (Kingma and Ba, 2015) with an initial learning rate of 1e-4. A mini-batch size of 32 subvolumes was used with data parallelization across 4 12GB NVIDIA Titan X Pascal GPUs was used for training and a mini-batch size of 8 subvolumes on 1 12GB NVIDIA Titan X Pascal GPU was used for validation and testing.

2.4. Quantifying Performance

2.4.1. Segmentation Performance Measure

To measure the quality of the produced segmentations, we calculated the Dice coefficient, which is defined by

$$Dice_c = \frac{2|\hat{y}_c \cdot \mathbf{y}_c|}{||\hat{y}_c||^2 + ||\mathbf{y}_c||^2} = \frac{2TP_c}{2TP_c + FN_c + FP_c}, \tag{19}$$

where \hat{y}_c is the binary segmentation for class c produced by a network, \mathbf{y}_c is the ground truth produced by FreeSurfer, TP_c is the true positive rate for class c , FN_c is the false negative rate for class c , and FP_c is the false positive rate for class c . We calculate the Dice coefficient separately for each class $c = 1, \dots, 50$, and average across classes to compute the overall performance of a network for one sMRI volume.

2.4.2. Uncertainty Measure

We quantify the uncertainty of a prediction, $p(y_m|\mathbf{x}_m)$, using the aleatoric uncertainty (Der Kiureghian and Ditlevsen, 2009; Kendall and Gal, 2017) for each voxel, \mathbf{v} . This was measured by the entropy of the softmax across the 50 output classes,

$$H(Y_{m,\mathbf{v}}|X_m = \mathbf{x}_m) = - \sum_{c=1}^{50} p(y_{m,\mathbf{v}} = c|\mathbf{x}_m) \log p(y_{m,\mathbf{v}} = c|\mathbf{x}_m), \tag{20}$$

We calculated the uncertainty for one sMRI volume by averaging across all output voxels not classified as background (i.e., given the unknown label).

3. RESULTS

3.1. Segmentation Performance

We trained MAP, MC Bernoulli Dropout (BD), and Spike-and-Slab Dropout (SSD) Meshnet-like CNNs on the 9,184 sMRI

TABLE 3 | The average and standard deviation of the class Dices across test volumes for the maximum a posteriori (MAP), MC Bernoulli dropout (BD), and spike-and-slab dropout (SSD) network on the in-site and out-of-site test sets.

Method	In-site	Out-of-site
MAP	0.7790 ± 0.0576	0.7333 ± 0.0498
BD	0.7764 ± 0.0506	0.7369 ± 0.0474
SSD	0.8373 ± 0.0471	0.7921 ± 0.0444

volumes in the training set. We then applied our networks to produce segmentations for both the in-site test set and the out-of-site test data. For the BD and SSD networks, 10 MC samples were used for test predictions. The means and standard deviations across volumes for the average Dice across all 50 classes are shown in **Table 3**. Dice scores for each label for the in-site and out-of-site test sets are shown in **Figures 2, 3**, respectively. We found that, compared to MAP and BD, SSD significantly increased the Dice for both the in-site ($p < 1e - 6$) and out-of-site ($p < 1e - 6$) test sets, per a paired t -test across test volumes. We found that SSD had a 5.7% drop in performance from the in-site test set to the out-of-site test set, where as the MAP has a drop of 6.2% and BD a drop of 5.4%. This is better than drops of 9.4% and 7.8% on average reported in the literature by Roy et al. (2018b) and Roy et al. (2018a), respectively. In **Figures 4, 5**, we show selected example segmentations for the SSD network for volumes that have Dice scores similar to the average Dice score across the respective dataset.

3.2. Utilizing Uncertainty

3.2.1. Predicting Segmentation Errors From Uncertainty

Ideally, an increase in DNN prediction uncertainty indicates an increase in the likelihood that prediction is incorrect. To evaluate whether this is the case for the trained brain segmentation DNN, we performed a receiver operating characteristic (ROC) analysis. In this analysis, voxels are ranked from most uncertain to least uncertain and one considers, at each rank, what fraction of the voxels were also misclassified by the network. An ROC curve can then be generated by plotting the true positive rate vs. the false negative rate for different uncertainty thresholds used to predict misclassification. The area under this curve (AUC) typically summarizes the results of the ROC analysis. The average ROC and AUCs across volumes for MAP, BD, and SSD for the in-site and out-of-site test sets are shown in **Figure 6**. Compared to MAP and BD, SSD significantly improved the AUC for both the in-site ($p < 1e - 6$) and out-of-site ($p < 1e - 6$) test sets, per a paired t -test across test set volumes.

3.2.2. Predicting Scan Quality From Uncertainty

Ideally, the output uncertainty for inputs not drawn from the training distribution should be relatively high. This could potentially be useful for a variety of applications. One particular application is detection of bad quality sMRI scans, since the

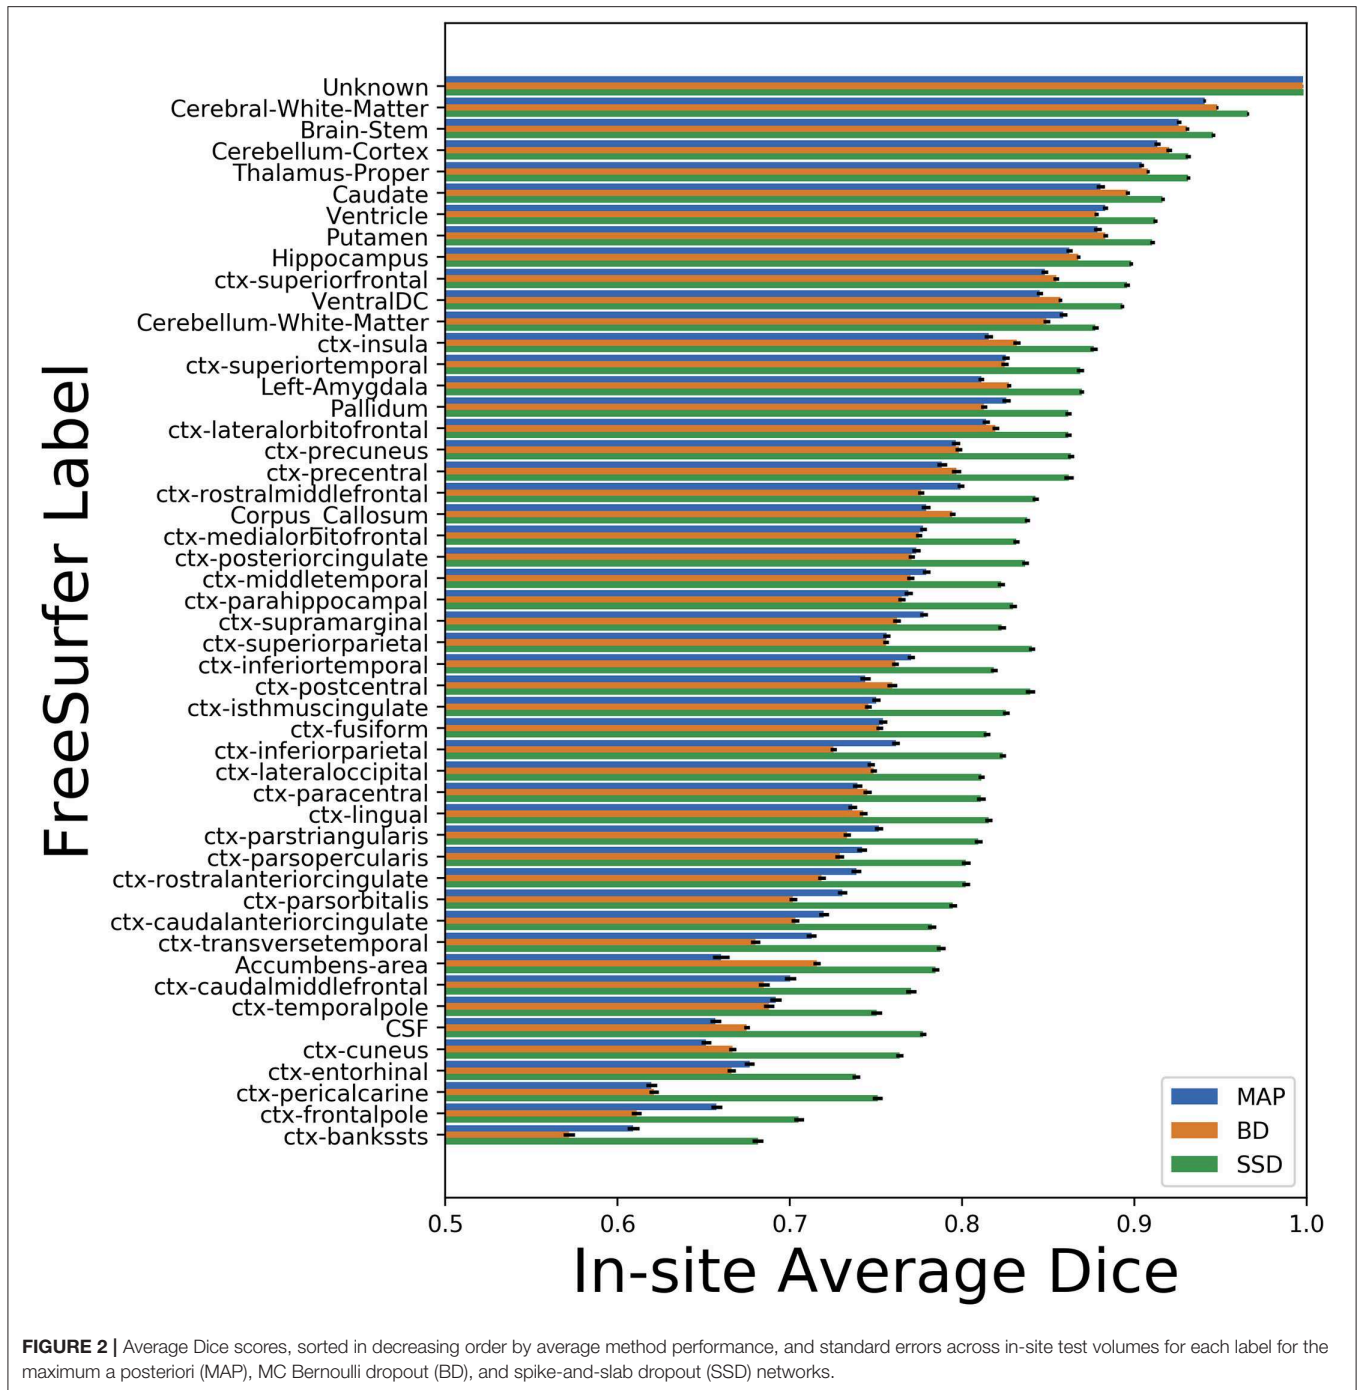


FIGURE 2 | Average Dice scores, sorted in decreasing order by average method performance, and standard errors across in-site test volumes for each label for the maximum a posteriori (MAP), MC Bernoulli dropout (BD), and spike-and-slab dropout (SSD) networks.

segmentation DNN was trained using relatively good quality scans. To test the validity of predicting high vs. low quality scans, we performed an ROC analysis on the held-out NNDSP dataset, where manual quality control ratings are available. We also did the same analysis using MRIQC (v0.10.5) (Esteban et al., 2017), a recently published method that combines a wide range of automated QC algorithms. To statistically test whether any method significantly outperformed the other methods, we performed bootstrap sampling of the AUC for predicting scan

quality from average uncertainty by sampling out-of-site test volumes. We performed 10,000 bootstrap samples, each with 418 volumes. The average ROC and AUC for the MAP, BD, SSD, and MRIQC methods are shown in **Figure 7**. The MAP, BD, and SSD networks all have significantly higher AUCs than MRIQC ($p = 1.369e - 4$, $p = 1.272e - 5$, and $p = 1.381e - 6$, respectively). Additionally, SSD had a significantly higher AUC than both MAP and BD ($p = 1.156e - 3$ and $p = 1.042e - 3$, respectively).




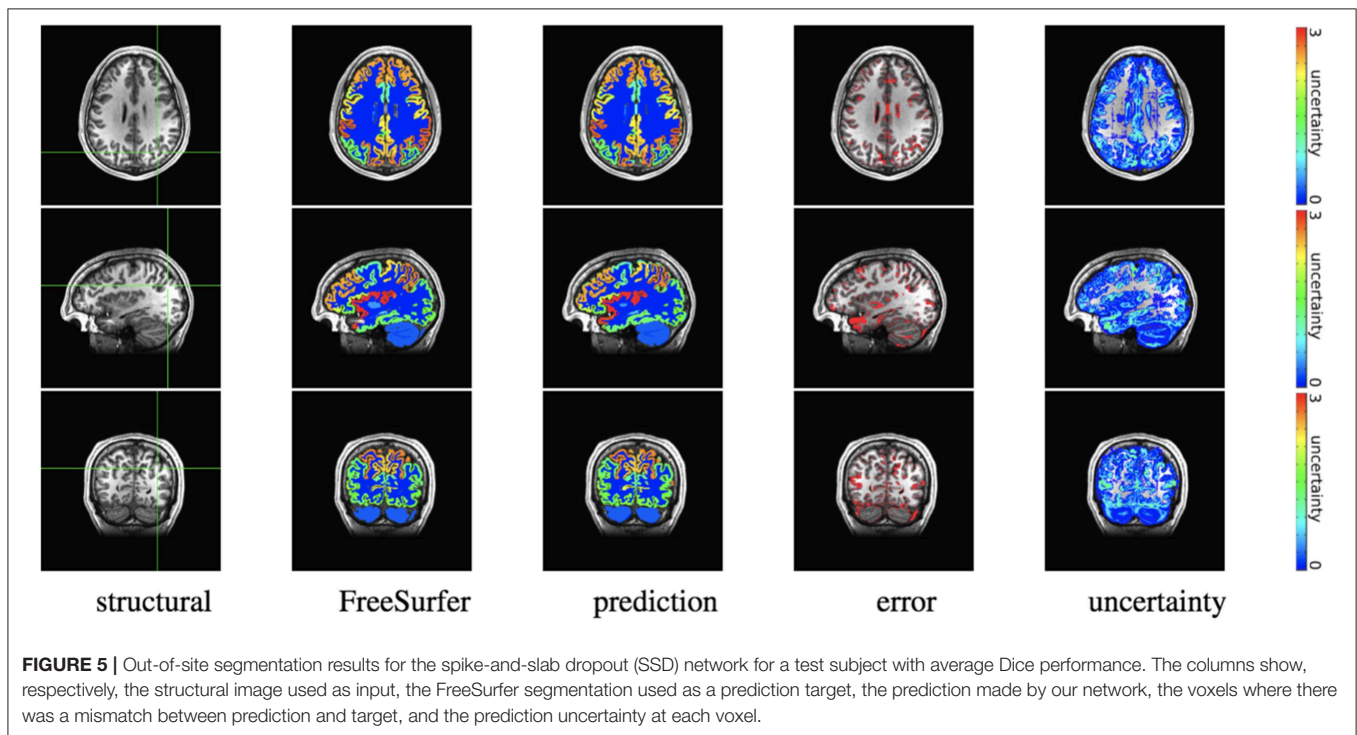
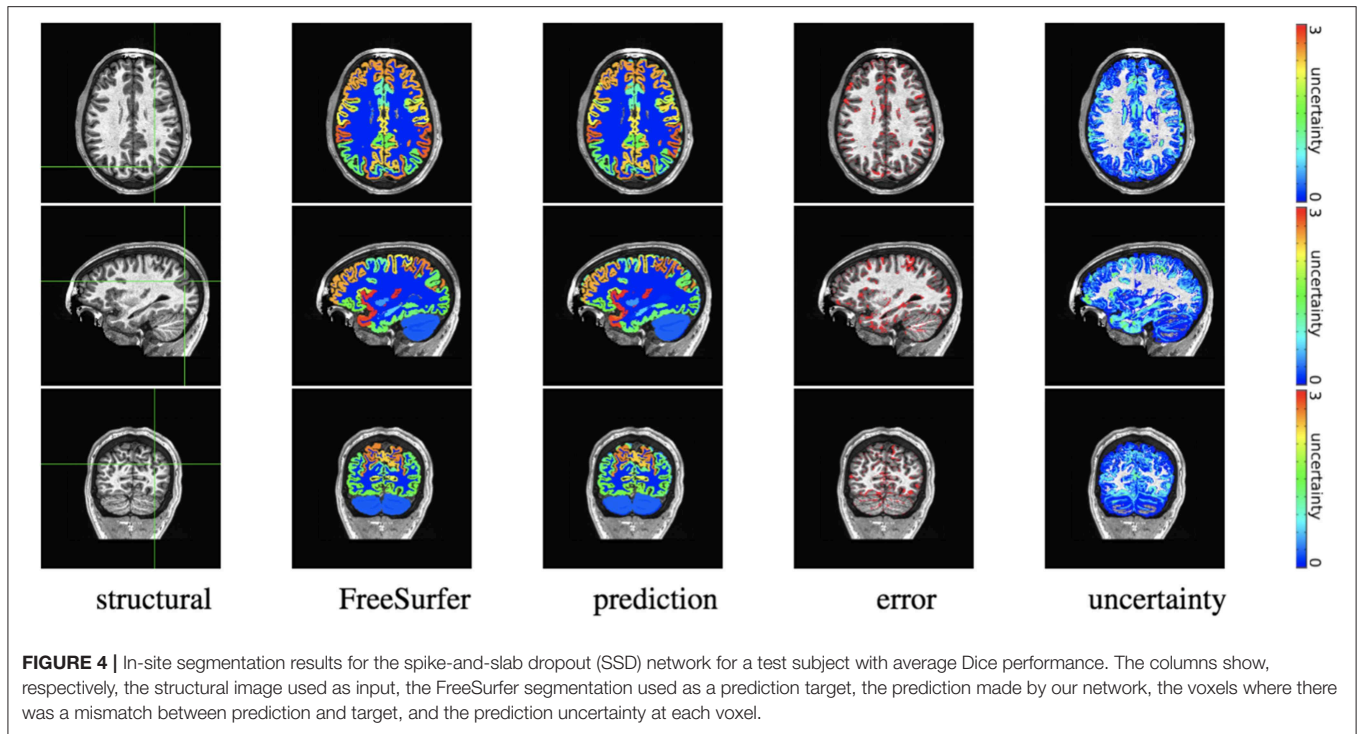
FIGURE 3 | Average Dice scores, sorted in decreasing order by average method performance, and standard errors across out-of-site test volumes for each label for the maximum a posteriori (MAP), MC Bernoulli dropout (BD), and spike-and-slab dropout (SSD) networks.

4. DISCUSSION

Segmentation of structures in sMRI volumes is a critical pre-processing step in many neuroimaging analyses. However, these segmentations are currently generated using tools that can take a day or more for each subject (FreeSurfer, 2018), such as FreeSurfer. This computational cost can be prohibitive when scaling analyses up from hundreds to thousands of subjects. DNNs have recently been proposed to perform sMRI segmentation in seconds to minutes. In this paper, we developed

a Bayesian DNN, using spike-and-slab dropout, with the goals of increasing the similarity of the DNN's predictions to the FreeSurfer segmentations and generating useful uncertainty estimates for these predictions.

In order to evaluate the proposed Bayesian network, we trained a standard deep neural network (DNN), using MAP estimation, to predict FreeSurfer segmentations from structural MRI (sMRI) volumes. We trained on a little under 10,000 sMRIs, obtained by combining approximately 70 different datasets (many of which, in turn, contain images from several sites,



e.g., NKI, ABIDE, ADHD200). We used a separate test set of more than 1,000 sMRIs, drawn from the same datasets. The resulting standard DNN performs at the same level of state-of-the-art networks (Fedorov et al., 2017a). This result, however, was obtained by testing over an order of magnitude more test data, and many more sites, than those papers. We also tested

performance on a completely separate dataset (NNDSP) from a site not encountered in training, which contained 418 sMRI volumes. Whereas Dice performance dropped slightly, this was less than what was observed in other studies (Roy et al., 2018a,b); this suggests that we may be achieving better generalization by training on our larger and more diverse dataset, and we plan on

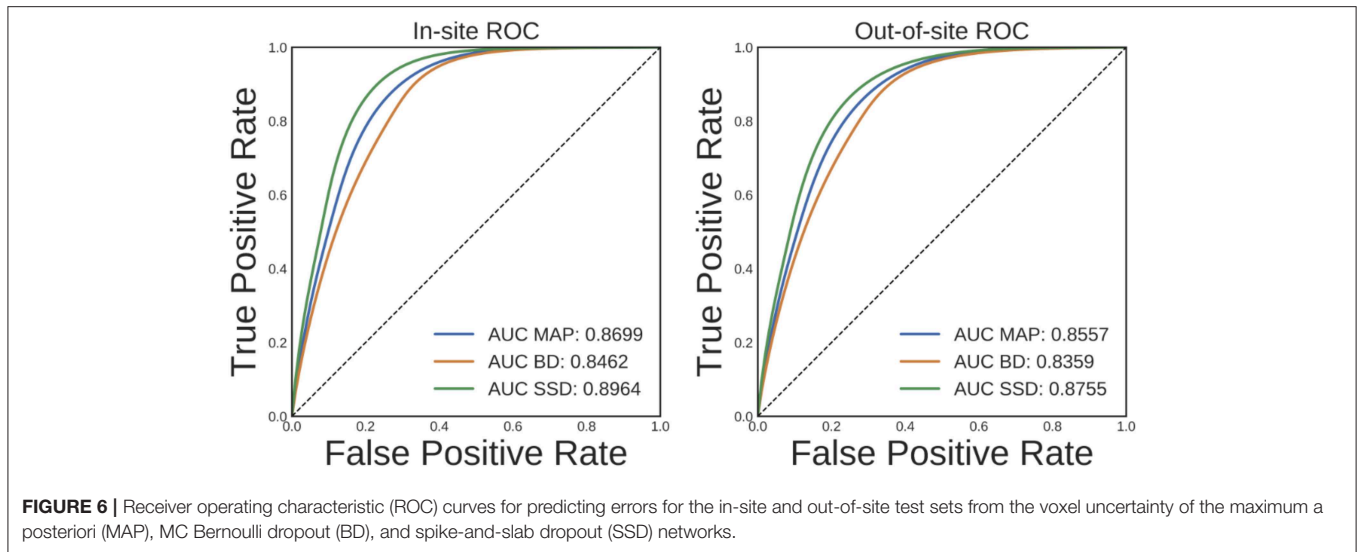


FIGURE 6 | Receiver operating characteristic (ROC) curves for predicting errors for the in-site and out-of-site test sets from the voxel uncertainty of the maximum a posteriori (MAP), MC Bernoulli dropout (BD), and spike-and-slab dropout (SSD) networks.




FIGURE 7 | Receiver operating characteristic (ROC) curves for predicting scan quality for the NNDSP out-of-site test set from the average non-background voxel uncertainty of the maximum a posteriori (MAP), MC Bernoulli dropout (BD), and spike-and-slab dropout (SSD) networks and from MRIQC scores.

testing this on more datasets from novel sites in the future. This is particularly important to us, as this network is meant to be used within an off-the-shelf tool¹.

We demonstrated that the estimated uncertainty for the prediction at each voxel is a good indicator of whether the standard network makes an error in it, both in-site and out-of-site. The tool that produces the predicted segmentation volume for an input sMRI will also produce an uncertainty volume. We anticipate this being useful at various levels, e.g., to refine other tools that rely on segmentation images, or to improve prediction models based on sMRI data (e.g., modification of calculation of

cortical thickness, surface area, voxel selection or weighting in regression (Roy et al., 2018a) or classification models, etc).

We also demonstrated that the average prediction uncertainty across voxels in the brain is an excellent indicator of manual quality control ratings. Furthermore, it outperforms the best existing automated solution (Esteban et al., 2017). Since automation is already used in large repositories (e.g., OpenMRI), we plan on offering our tool as an additional quality control measure.

Finally, we showed that a new Bayesian DNN using spike-and-slab dropout with learned model uncertainty was significantly better than previous approaches. This spike-and-slab method increased segmentation performance and improved the usefulness of output uncertainties compared both to a MAP DNN method and an MC Bernoulli dropout method, which has previously been used in the brain segmentation literature (Li et al., 2017; Roy et al., 2018a). These results show that Bayesian DNNs are a promising method for building brain segmentation and automated sMRI quality control tools. We have also made a version of “Nobrainer,” that incorporates the networks trained and evaluated in this paper, available for download and use within a Singularity/Docker container².

We believe it may be possible to improve this segmentation processing, in that we did not use registration. One option would be to use various techniques for data augmentation (e.g., variation of image contrast, since that is pretty heterogeneous, rotations/translations of existing examples, addition of realistic noise, etc). Another would be to eliminate the need to divide the brain into sub-volumes, which loses some global information; this will become more feasible in GPUs with more memory. Finally, we plan on using post-processing of results (e.g., ensure some coherence between predictions for adjacent voxels, leverage off-the-shelf brain and tissue masking code).

²<https://github.com/neuronets/kwyk>

DATA AVAILABILITY STATEMENT

All datasets generated for this study are included in the manuscript/supplementary files.

AUTHOR CONTRIBUTIONS

PM and CZ developed the novel Bayesian neural network method. PM, NR, JK, and SG wrote the code for training and testing the neural networks. PM and NR analyzed the neural network predictions. JL, DN, and AT acquired and preprocessed the magnetic

resonance imaging (MRI) data. PM and FP wrote the first manuscript draft. All authors contributed to manuscript revision, read the manuscript, and approved the submitted version.

ACKNOWLEDGMENTS

This research was supported (in part) by the Intramural Research Program of the NIMH (ZICMH002968). This work utilized the computational resources of the NIH HPC Biowulf cluster (<http://hpc.nih.gov>). JK's and SG's contribution was supported by NIH R01 EB020740.

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). "Tensorflow: a system for large-scale machine learning," in *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)* (Savannah, GA), 265–283.
- Alexander, L. M., Escalera, J., Ai, L., Andreotti, C., Febre, K., Mangone, A., et al. (2017). An open resource for transdiagnostic research in pediatric mental health and learning disorders. *Sci. Data* 4:170181. doi: 10.1038/sdata.2017.181
- Bellec, P., Chu, C., Chouinard-Decorte, F., Benhajali, Y., Margulies, D. S., and Craddock, R. C. (2017). The neuro bureau adhd-200 preprocessed repository. *Neuroimage* 144, 275–286. doi: 10.1016/j.neuroimage.2016.06.034
- Biswal, B. B., Mennes, M., Zuo, X.-N., Gohel, S., Kelly, C., Smith, S. M., et al. (2010). Toward discovery science of human brain function. *Proc. Natl. Acad. Sci. U.S.A.* 107, 4734–4739. doi: 10.1073/pnas.0911855107
- Blumenthal, J. D., Zijdenbos, A., Molloy, E., and Giedd, J. N. (2002). Motion artifact in magnetic resonance imaging: implications for automated analysis. *Neuroimage* 16, 89–92. doi: 10.1006/nimg.2002.1076
- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. (2015). "Weight uncertainty in neural networks," in *International Conference on Machine Learning (Lille)*, 1613–1622.
- Cardoso, M. J., Modat, M., Wolz, R., Melbourne, A., Cash, D., Rueckert, D., et al. (2015). Geodesic information flows: spatially-variant graphs and their application to segmentation and fusion. *IEEE Trans. Med. Imaging* 34, 1976–1988. doi: 10.1109/TMI.2015.2418298
- Der Kiureghian, A., and Ditlevsen, O. (2009). Aleatory or epistemic? Does it matter? *Struct. Safety* 31, 105–112. doi: 10.1016/j.strusafe.2008.06.020
- Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F. X., Alaerts, K., et al. (2014). The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol. Psychiatry* 19:659. doi: 10.1038/mp.2013.78
- di Oleggio Castello, M. V., Halchenko, Y. O., Guntupalli, J. S., Gors, J. D., and Gobbini, M. I. (2017). The neural representation of personally familiar and unfamiliar faces in the distributed system for face perception. *Sci. Rep.* 7. doi: 10.1038/s41598-017-12559-1
- Dolz, J., Desrosiers, C., and Ayed, I. B. (2018). 3d fully convolutional networks for subcortical segmentation in MRI: a large-scale study. *NeuroImage* 170, 456–470. doi: 10.1016/j.neuroimage.2017.04.039
- Esteban, O., Birman, D., Schaer, M., Koyejo, O. O., Poldrack, R. A., and Gorgolewski, K. J. (2017). Mriqc: advancing the automatic prediction of image quality in mri from unseen sites. *PLoS ONE* 12:e0184661. doi: 10.1371/journal.pone.0184661
- Fedorov, A., Damaraju, E., Calhoun, V., and Plis, S. (2017a). Almost instant brain atlas segmentation for large-scale studies. *arXiv:1711.00457*.
- Fedorov, A., Johnson, J., Damaraju, E., Ozerin, A., Calhoun, V., and Plis, S. (2017b). "End-to-end learning of brain tissue segmentation from imperfect labeling," in *International Joint Conference on Neural Networks (Anchorage, AK: IEEE)*, 3785–3792.
- Fischl, B. (2012). Freesurfer. *Neuroimage* 62. doi: 10.1016/j.neuroimage.2012.01.021
- FreeSurfer (2018). *Reconallruntimes*. (accessed December 4, 2018).
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., and Frackowiak, R. S. J. (1994). Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210.
- Gal, Y. (2016). *Uncertainty in deep learning*. (Ph.D. thesis). Cambridge, UK: University of Cambridge.
- Gal, Y., and Ghahramani, Z. (2016). "Dropout as a Bayesian approximation: representing model uncertainty in deep learning," in *International Conference on Machine Learning (New York, NY)*, 1050–1059.
- Gal, Y., Hron, J., and Kendall, A. (2017). "Concrete dropout," in *Advances in Neural Information Processing Systems (Long Beach, CA)*, 3581–3590.
- Graves, A. (2011). "Practical variational inference for neural networks," in *Advances in Neural Information Processing Systems (Granda)*, 2348–2356.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). "On calibration of modern neural networks," in *Proceedings of the 34th International Conference on Machine Learning, Vol. 70 (Sydney, NSW)*, 1321–1330.
- Hastie, T., Tibshirani, R., Friedman, J., and Franklin, J. (2005). The elements of statistical learning: data mining, inference and prediction. *Math. Intell.* 27, 83–85. doi: 10.1007/BF02985802
- Haxby, J. V., Guntupalli, J. S., Connolly, A. C., Halchenko, Y. O., Conroy, B. R., Gobbini, M. I., et al. (2011). A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* 72, 404–416. doi: 10.1016/j.neuron.2011.08.026
- Hinton, G. E., and Van Camp, D. (1993). "Keeping the neural networks simple by minimizing the description length of the weights," in *Proceedings of the sixth annual conference on Computational learning theory (Santa Cruz, CA: ACM)*, 5–13.
- Holmes, A. J., Hollinshead, M. O., OKeefe, T. M., Petrov, V. I., Fariello, G. R., Wald, L. L., et al. (2015). Brain genomics superstruct project initial data release with structural, functional, and behavioral measures. *Sci. Data* 2:150031. doi: 10.1038/sdata.2015.31
- Hron, J., Matthews, A. G., and Ghahramani, Z. (2018). "Variational bayesian dropout: pitfalls and fixes," in *International Conference on Machine Learning (Stockholm)*, 2024–2033.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., and Smith, S. M. (2012). FSL. *Neuroimage* 62, 782–790. doi: 10.1016/j.neuroimage.2011.09.015
- Keator, D. B., van Erp, T. G., Turner, J. A., Glover, G. H., Mueller, B. A., Liu, T. T., et al. (2016). The function biomedical informatics research network data repository. *Neuroimage* 124, 1074–1079. doi: 10.1016/j.neuroimage.2015.09.003
- Kendall, A., and Gal, Y. (2017). "What uncertainties do we need in bayesian deep learning for computer vision?," in *Advances in Neural Information Processing Systems (Long Beach, CA)*, 5574–5584.
- Kennedy, D. N., Haselgrove, C., Hodge, S. M., Rane, P. S., Makris, N., and Frazier, J. A. (2012). CANDIShare: a resource for pediatric neuroimaging data. *Neuroinformatics*. 10, 319–322. doi: 10.1007/s12021-011-9133-y
- Keshavan, A., Yeatman, J., and Rokem, A. (2018). Combining citizen science and deep learning to amplify expertise in neuroimaging. *bioRxiv*. doi: 10.1101/363382

- Kingma, D. P., and Ba, J. (2015). "Adam: a method for stochastic optimization," in *International Conference on Learning Representations* (San Diego, CA).
- Kingma, D. P., Salimans, T., and Welling, M. (2015). "Variational dropout and the local reparameterization trick," in *Advances in Neural Information Processing Systems* (Montreal, QC), 2575–2583.
- Ledig, C., Heckemann, R. A., Hammers, A., Lopez, J. C., Newcombe, V. F., Makropoulos, A., et al. (2015). Robust whole-brain segmentation: application to traumatic brain injury. *Med. Image Anal.* 21, 40–58. doi: 10.1016/j.media.2014.12.003
- Lee, J. A., Gustavo, S., Migineishvili, N., Nielson, D., Bandettini, P. A., Shaw, P., et al. (2018). "Automated quality control methods applied to a novel dataset," in *Organization for Human Brain Mapping (OHBM) Annual Meeting* (Singapore).
- Li, W., Wang, G., Fidon, L., Ourselin, S., Cardoso, M. J., and Vercauteren, T. (2017). "On the compactness, efficiency, and representation of 3D convolutional networks: brain parcellation as a pretext task," in *International Conference on Information Processing in Medical Imaging* (Boone, NC: Springer), 348–360.
- Liu, W., Wei, D., Chen, Q., Yang, W., Meng, J., Wu, G., et al. (2017). Longitudinal test-retest neuroimaging data from healthy young adults in Southwest China. *Sci. Data* 4:170017. doi: 10.1038/sdata.2017.17
- Louizos, C., and Welling, M. (2017). "Multiplicative normalizing flows for variational Bayesian neural networks," in *International Conference on Machine Learning* (Sydney, NSW), 2218–2227.
- Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., and Buckner, R. L. (2007). Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* 19, 1498–1507. doi: 10.1162/jocn.2007.19.9.1498
- Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., et al. (2001). A probabilistic atlas and reference system for the human brain: international consortium for brain mapping (icbm). *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 356, 1293–1322. doi: 10.1098/rstb.2001.0915
- McClure, P., and Kriegeskorte, N. (2017). "Robustly representing uncertainty in deep neural networks through sampling," in *NIPS Bayesian Deep Learning Workshop* (Long Beach, CA).
- McClure, P., Zheng, C. Y., Kaczmarzyk, J., Rogers-Lee, J., Ghosh, S., Nielson, D., et al. (2018). "Distributed weight consolidation: a brain segmentation case study," in *Advances in Neural Information Processing Systems 31* (Montreal, QC), 4093–4103.
- Molchanov, D., Ashukha, A., and Vetrov, D. (2017). "Variational dropout sparsifies deep neural networks," in *International Conference on Machine Learning* (Sydney, NSW), 2498–2507.
- Mueller, S. G., Weiner, M. W., Thal, L. J., Petersen, R. C., Jack, C., Jagust, W., et al. (2005). The alzheimer's disease neuroimaging initiative. *Neuroimaging Clin.* 15, 869–877. doi: 10.1016/j.nic.2005.09.008
- Nastase, S. A., Connolly, A. C., Oosterhof, N. N., Halchenko, Y. O., Guntupalli, J. S., Visconti di Oleggio Castello, M., et al. (2017). Attention selectively reshapes the geometry of distributed semantic representation. *Cereb. Cortex* 27, 4277–4291. doi: 10.1093/cercor/bhx138
- Nguyen, C. V., Li, Y., Bui, T. D., and Turner, R. E. (2018). "Variational continual learning," in *International Conference on Learning Representations* (Vancouver, BC).
- Nooner, K. B., Colcombe, S., Tobe, R., Mennes, M., Benedict, M., Moreno, A., et al. (2012). The NKI-Rockland sample: a model for accelerating the pace of discovery science in psychiatry. *Front. Neurosci.* 6:152. doi: 10.3389/fnins.2012.00152
- O'Connor, D., Potler, N. V., Kovacs, M., Xu, T., Ai, L., Pellman, J., et al. (2017). The healthy brain network serial scanning initiative: a resource for evaluating inter-individual differences and their reliabilities across scan conditions and sessions. *Gigascience* 6, 1–14. doi: 10.1093/gigascience/giw011
- Petersen, R. C., Aisen, P., Beckett, L., Donohue, M., Gamst, A., Harvey, D., et al. (2010). Alzheimer's disease neuroimaging initiative (ADNI): clinical characterization. *Neurology* 74, 201–209. doi: 10.1212/WNL.0b013e3181c6b3e25
- Poldrack, R. A., Barch, D. M., Mitchell, J., Wager, T., Wagner, A. D., Devlin, J. T., et al. (2013). Toward open sharing of task-based fmri data: the openfmri project. *Front. Neuroinform.* 7:12. doi: 10.3389/fninf.2013.00012
- Rajchl, M., Pawlowski, N., Rueckert, D., Matthews, P. M., and Glocker, B. (2018). "Neuronet: fast and robust reproduction of multiple brain image segmentation pipelines," in *Medical Imaging With Deep Learning* (Amsterdam).
- Robbins, H., and Monro, S. (1951). A stochastic approximation method. *Ann. Math. Stat.* 22, 400–407. doi: 10.1214/aoms/1177729586
- Rohlfing, T. (2012). Image similarity and tissue overlaps as surrogates for image registration accuracy: widely used but unreliable. *IEEE Trans. Med. Imaging* 31, 153–163. doi: 10.1109/TMI.2011.2163944
- Ronneberger, O., Fischer, P., and Brox, T. (2015). "U-net: convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Munich: Springer), 234–241.
- Roy, A. G., Conjeti, S., Navab, N., and Wachinger, C. (2018a). Bayesian quicknat: model uncertainty in deep whole-brain segmentation for structure-wise quality control. *arXiv:1811.09800*. doi: 10.1016/j.neuroimage.2019.03.042
- Roy, A. G., Conjeti, S., Navab, N., and Wachinger, C. (2018b). QuickNAT: segmenting MRI neuroanatomy in 20 seconds. *arXiv:1801.04161*.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.
- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., et al. (2015). UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* 12:e1001779. doi: 10.1371/journal.pmed.1001779
- Titsias, M. K., and Lázaro-Gredilla, M. (2011). "Spike and slab variational inference for multi-task and multiple kernel learning," in *Advances in Neural Information Processing Systems* (Granada), 2339–2347.
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., et al. (2013). The WU-Minn human connectome project: an overview. *NeuroImage* 80, 62–79. doi: 10.1016/j.neuroimage.2013.05.041
- Vzquez, P. G., Whitfield-Gabrieli, S., Bauer, C. C., Barrios, and A. F. (2016). Brain functional connectivity of hypnosis without target suggestion. an intrinsic hypnosis rs-fMRI study.
- Wei, D., Zhuang, K., Chen, Q., Yang, W., Liu, W., Wang, K., et al. (2018). Structural and functional MRI from a cross-sectional Southwest University adult lifespan dataset (sald). *bioRxiv* 177279. doi: 10.1101/177279
- Yu, F., and Koltun, V. (2015). "Multi-scale context aggregation by dilated convolutions," in *International Conference on Learning Representations* (San Diego, CA).
- Zuo, X.-N., Anderson, J. S., Bellec, P., Birn, R. M., Biswal, B. B., Blautzik, J., et al. (2014). An open science resource for establishing reliability and reproducibility in functional connectomics. *Sci. Data* 1:140049.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

At least a portion of this work is authored by Patrick McClure, Nao Rho, John A. Lee, Charles Y. Zheng, Dylan M. Nielson, Adam G. Thomas, Peter Bandettini, and Francisco Pereira on behalf of the U.S. Government and, as regards Dr. McClure, Mr. Rho, Dr. Lee, Dr. Zheng, Dr. Nielson, Dr. Thomas, Dr. Bandettini, Dr. Pereira, and the U.S. Government, is not subject to copyright protection in the United States. Foreign and other copyrights may apply. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution, or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.