

Fully Bayesian VIB-DeepSSM

Jadie Adams^{1,2} and Shireen Elhabian^{1,2}

¹ Scientific Computing and Imaging Institute, University of Utah, UT, USA

² School of Computing, University of Utah, UT, USA
 jadie.adams@utah.edu, shireen@sci.utah.edu

Abstract. Statistical shape modeling (SSM) enables population-based quantitative analysis of anatomical shapes, informing clinical diagnosis. Deep learning approaches predict correspondence-based SSM directly from unsegmented 3D images but require calibrated uncertainty quantification, motivating Bayesian formulations. Variational information bottleneck DeepSSM (VIB-DeepSSM) is an effective, principled framework for predicting probabilistic shapes of anatomy from images with aleatoric uncertainty quantification. However, VIB is only half-Bayesian and lacks epistemic uncertainty inference. We derive a fully Bayesian VIB formulation from both the probably approximately correct (PAC)-Bayes and variational inference perspectives. We demonstrate the efficacy of two scalable approaches for Bayesian VIB with epistemic uncertainty: concrete dropout and batch ensemble. Additionally, we introduce a novel combination of the two that further enhances uncertainty calibration via multimodal marginalization. Experiments on synthetic shapes and left atrium data demonstrate that the fully Bayesian VIB network predicts SSM from images with improved uncertainty reasoning without sacrificing accuracy.

Keywords: Statistical Shape Modeling, Bayesian Deep Learning, Variational Information Bottleneck, Epistemic Uncertainty Quantification

1 Introduction

Statistical Shape Modeling (SSM) is a powerful tool for describing anatomical shapes (i.e., bones and organs) in relation to a cohort of interest. Correspondence-based shape modeling is popular due to its interpretable shape representation using landmarks or correspondence points on anatomical surfaces that are spatially consistent across the population. Specifically, each shape is represented by a dense set of correspondences, denoted as a point distribution model (PDM), that is automatically defined on shapes (e.g., via optimization [8] or pairwise parameterization [23]) segmented from 3D medical images. Conventional SSM pipelines require expert-driven, intensive steps such as segmentation, shape registration, and tuning correspondence optimization parameters or defining an atlas/template for pairwise surface matching. Deep learning approaches have mitigated this overhead by providing end-to-end solutions, predicting PDMs from unsegmented 3D images with little preprocessing [5,1,25,2]. Such solutions cannot be safely deployed in sensitive, clinical decision-making scenarios without

uncertainty reasoning [11], which provides necessary insight into the degree of model confidence and serves as a metric of prediction reliability. There are two primary forms of uncertainty aleatoric (or data-dependent) and epistemic (or model-dependent) [17]. The overall prediction uncertainty is a sum of the two. It is essential to distinguish between these two forms, as epistemic is reducible and can be decreased given more training data or by refining the model [10]. Bayesian deep learning frameworks automatically provide epistemic uncertainty reasoning and can be defined to predict distributions, providing aleatoric uncertainty quantification [7,17,19].

DeepSSM [5] is a state-of-the-art framework providing SSM estimates that perform statistically similarly to traditional SSM methods in downstream tasks [6]. Uncertain-DeepSSM [1] adapted the DeepSSM network to be Bayesian, providing both forms of uncertainties. DeepSSM, Uncertain-DeepSSM, and other formulations [25] rely on a shape prior in the form of a supervised latent encoding pre-computed using principal component analysis (PCA). PCA supervision imposes a linear relationship between the latent and the output space, restricts the learning task, and does not scale in the case of large sets of high-dimensional shape data. VIB-DeepSSM [2] relaxes these assumptions to provide improved accuracy and aleatoric uncertainty estimates over the existing state-of-the-art methods [5,1,25]. This probabilistic formulation utilizes a variational information bottleneck (VIB) [3] architecture to learn the latent encoding in the context of the task. VIB-DeepSSM is self-regularized via a latent prior, increasing generalizability and helping alleviate the need for the computationally expensive DeepSSM data augmentation process. However, this approach does not quantify epistemic uncertainty because VIB is only half-Bayesian [3]. In this paper, we propose to significantly extend the VIB-DeepSSM framework to be fully Bayesian, predicting probabilistic anatomy shapes directly from images with both forms of uncertainty quantification.

The contributions of this work include the following: (1) We mathematically derive fully Bayesian VIB from two perspectives: PAC-Bayes bounds and variational inference. (2) We demonstrate two scalable approaches for Bayesian VIB-DeepSSM with epistemic uncertainty quantification (concrete dropout and batch ensemble) and compare them to naive ensembling. (3) We introduce and theoretically justify a novel combination of concrete dropout and ensembling for improved uncertainty calibration. (4) We illustrate that the fully Bayesian formulations improve uncertainty reasoning (especially the proposed method) on synthetic and real data without sacrificing accuracy.

2 Background

We denote a set of paired training data as $\mathcal{D} = \{\mathcal{X}, \mathcal{Y}\}$. $\mathcal{X} = \{\mathbf{x}_n\}_{n=1}^N$ is a set of N unsegmented images, where $\mathbf{x}_n \in \mathbb{R}^{H \times W \times D}$. $\mathcal{Y} = \{\mathbf{y}_n\}_{n=1}^N$ is the set of PDMs comprised of M 3D correspondence points, where $\mathbf{y}_n \in \mathbb{R}^{3M}$. VIB utilizes a stochastic latent encoding $\mathcal{Z} = \{\mathbf{z}_n\}_{n=1}^N$, where $\mathbf{z}_n \in \mathbb{R}^L$ and $L \ll 3M$.

In **Bayesian modeling**, model parameters Θ are obtained by maximizing the likelihood $p(\mathbf{y}|\mathbf{x}, \Theta)$. The predictive distribution is found by marginalizing over Θ , which requires solving for the posterior $p(\Theta|\mathcal{D})$. In most cases, $p(\Theta|\mathcal{D})$ is not analytically tractable; thus, an approximate posterior $q(\Theta)$ is found via variational inference (VI). Bayesian networks maximize the VI evidence lower bound (ELBO) by minimizing:

$$v_I = \mathbb{E}_{\tilde{\Theta} \sim q(\Theta)} \left[-\log p(\mathbf{y}|\mathbf{x}, \tilde{\Theta}) \right] + \beta \text{KL} [q(\Theta) \| p(\Theta)] \quad (1)$$

where $p(\Theta)$ is the prior on network weights, and β is a weighting parameter.

The deep **Variational Information Bottleneck** (VIB) [3] model learns to predict \mathbf{y} from \mathbf{x} using a low dimensional stochastic encoding \mathbf{z} . The VIB architecture comprises of a stochastic encoder parameterized by ϕ , $q(\mathbf{z}|\mathbf{x}, \phi)$, and a decoder parameterized by θ , $p(\mathbf{y}|\mathbf{z}, \theta)$ (Figure 1). VIB utilizes VI to derive a theoretical lower bound on the information bottleneck objective:

$$v_{IB} = \mathbb{E}_{\hat{\mathbf{z}} \sim q(\mathbf{z}|\mathbf{x}, \phi)} \left[-\log p(\mathbf{y}|\hat{\mathbf{z}}, \theta) \right] + \beta \text{KL} [q(\mathbf{z}|\mathbf{x}, \phi) \| p(\mathbf{z})] \quad (2)$$

The entropy of the $p(\mathbf{y}|\mathbf{z})$ distribution (computed using sampling) captures aleatoric uncertainty. The VIB objective has also been derived using an alternative motivation: Bayesian inference via optimizing a PAC style upper bound on the true negative log-likelihood risk [4]. Through this PAC-Bayes lens, it has been proven that VIB is *half Bayesian*, as the Bayesian strategy is applied to minimize an upper bound with respect to the conditional expectation of \mathbf{y} , but the Maximum Likelihood Estimation (MLE) strategy is used to approximate the expectation over inputs. The VIB objective can be made a fully valid bound on the true risk by applying an additional PAC-Bound with respect to the parameters, resulting in a fully Bayesian VIB that captures epistemic uncertainty in addition to aleatoric.

3 Methods

3.1 Bayesian Variational Information Bottleneck

In fully-Bayesian VIB (BVIB), rather than fitting the model parameters $\Theta = \{\phi, \theta\}$ via MLE, we use VI to approximate the posterior $p(\Theta|\mathcal{D})$. There are now two intractable posteriors $p(\mathbf{z}|\mathbf{x}, \phi)$ and $p(\Theta|\mathbf{x}, \mathbf{y})$. The first is approximated via $q(\mathbf{z}|\mathbf{x}, \phi)$ as in Eq. 2 and the second is approximated by $q(\Theta)$ as in Eq. 1. Minimizing these two KL divergences via a joint ELBO gives the objective (see Appendix A for derivation details):

$$v_{\text{BVIB}} = \mathbb{E}_{\tilde{\Theta}} \left[\mathbb{E}_{\hat{\mathbf{z}}} \left[-\log p(\mathbf{y}|\hat{\mathbf{z}}, \tilde{\theta}) \right] + \text{KL} \left[q(\mathbf{z}|\mathbf{x}, \tilde{\phi}) \| p(\mathbf{z}) \right] \right] + \text{KL} [q(\Theta) \| p(\Theta)] \quad (3)$$

where $\tilde{\Theta} \sim q(\Theta)$ and $\hat{\mathbf{z}} \sim q(\mathbf{z}|\mathbf{x}, \tilde{\phi})$. This objective is equivalent to the BVIB objective acquired via applying a PAC-Bound with respect to the conditional expectation of targets and then another with respect to parameters [4]. This

is expected, as it has been proven that the VI formulation using ELBO and the PAC-Bayes formulation with negative log-likelihood as the risk metric are algorithmically identical [24]. Additionally, this matches the objective derived for the Bayesian VAE when $\mathbf{y} = \mathbf{x}$ [9]. Implementing BVIB requires defining a prior distribution for the latent representation $p(\mathbf{z})$ and the network weights $p(\Theta)$. Following VIB, we define, $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\mathbf{0}, \mathbb{I})$. Different methods exist for defining $p(\Theta)$, and multiple approaches are explored in the following section.

3.2 Proposed BVIB-DeepSSM Model Variants

In adapting VIB-DeepSSM to be fully Bayesian, we propose utilizing two approaches that have demonstratively captured epistemic uncertainty without greatly increasing computational and memory costs: concrete dropout [13] and batch ensemble [26]. Additionally, we propose a novel integration for a more flexible, multimodal posterior approximation.

Concrete Dropout (CD) utilizes Monte Carlo dropout sampling as a scalable solution for approximate VI [13]. Epistemic uncertainty is captured by the spread of predictions with sampled dropout masks in inference. CD automatically optimizes layer-wise dropout probabilities along with the network weights.

Naive Ensemble (NE) models combine outputs from several networks for improved performance. Networks trained with different initialization converge to different local minima, resulting in test prediction disagreement [12]. The spread in predictions effectively captures epistemic uncertainty [21]. NE models are computationally expensive, as cost increases linearly with number of members.

Batch Ensemble (BE) [26] compromises between a single network and NE, balancing the trade-off between accuracy and running time and memory. In BE, each weight matrix is defined to be the Hadamard product of a shared weight among all ensemble members and a rank-one matrix per member. BE provides an ensemble from one network, where the only extra computation cost is the Hadamard product, and the only added memory overhead is sets of 1D vectors.

Novel Integration of Dropout and Ensembling: Deep ensembles have historically been presented as a non-Bayesian competitor for uncertainty estimation. However, recent work argues that ensembles often approximate the predictive distribution more closely than canonical approximate inference procedures (i.e., VI) and are an effective mechanism for approximate Bayesian marginalization [28]. Furthermore, it has been shown that combining traditional Bayesian methods with ensembling improves the fidelity of approximate inference via multimodal marginalization, resulting in a more robust, accurate model [27]. In concrete dropout, the approximate variational distribution is parameterized via a concrete distribution. While this parameterization enables efficient Bayesian inference, it greatly limits the expressivity of the approximate posterior. To help remedy this, we propose integrating concrete dropout and ensembling (BE-CD and NE-CD) to acquire a multimodal approximate posterior on weights for increased flexibility and expressiveness. To the best of our knowledge, this combination has not previously been proposed with the motivation of multimodal marginalization for improved uncertainty calibration.




Fig. 1: Common VIB-DeepSSM Architecture for all proposed variants.

3.3 BVIB-DeepSSM Implementation

We compare the proposed BVIB approaches with the original VIB-DeepSSM formulation [2]. All models have the overall structure shown in Figure 1, comprised of a 3D convolutional encoder (f_e) and fully connected decoder (f_d). CD models have concrete dropout following every layer, BE weights have four members (the maximum GPU memory would allow), and four models were used to create NE models for a fair comparison. Following [2], burn-in is used to convert the loss from deterministic (L2) to probabilistic (Eqs 10, 3, 13). This counteracts the typical reduction in accuracy that occurs when a negative log-likelihood based loss is used with a gradient-based optimizer [22]. An additional dropout burn-in phase is used for CD models to increase the speed of convergence. All models were trained until the validation accuracy had not decreased in 50 epochs. A table of model hyperparameters and tested ranges is proved in Appendix C. The training was done on Tesla V100 GPU with Xavier initialization [15], Adam optimization [18]. The prediction uncertainty is a sum of the epistemic (variance resulting from marginalizing over Θ) and aleatoric (variance resulting from marginalizing over \mathbf{z}) uncertainty (see Appendix B for calculation details).

4 Results

We expect well-calibrated prediction uncertainty to correlate with the error, aleatoric uncertainty to correlate with the input image outlier degree (given that it is data-dependent), and epistemic uncertainty to correlate with the shape outlier degree (i.e., to detect out-of-distribution data). The outlier degree value for each mesh and image is quantified by running PCA (preserving 95% of variability) and then considering the Mahalanobis distance of the PCA scores to the mean (within-subspace distance) and the reconstruction error (off-subspace distance). The sum of these values provides a measure of similarity to the whole set in standard deviation units [20]. Experiments are designed to evaluate this expected correlation as well as accuracy, which is calculated as the root mean square error (RMSE) between the true and predicted points. Additionally, we quantify the *surface-to-surface distance* between a mesh reconstructed from the predicted PDM (predicted mesh) and the ground truth segmented mesh. The reported results are an average of four runs for each model, excluding the NE models, which ensemble the four runs.

4.1 Supershapes Experiments

Supershapes (SS) are synthetic 3D shapes parameterized by variables that determine the curvature and number of lobes [14]. We generated 1200 supershapes with lobes randomly selected between 3 and 7 and curvature parameters randomly sampled from a χ^2 distribution with 4 degrees of freedom. Corresponding 3D images were generated with foreground and background intensity values modeled as Gaussian distributions with different means and equal variance. Images were blurred with a Gaussian filter (size randomly selected between 1 and 8) to mimic diffuse shape boundaries. Figure 2A displays example shape meshes and images with corresponding outlier degrees, demonstrating the wide variation. We randomly split the mesh/image pairs to create a training set of size 1000, a

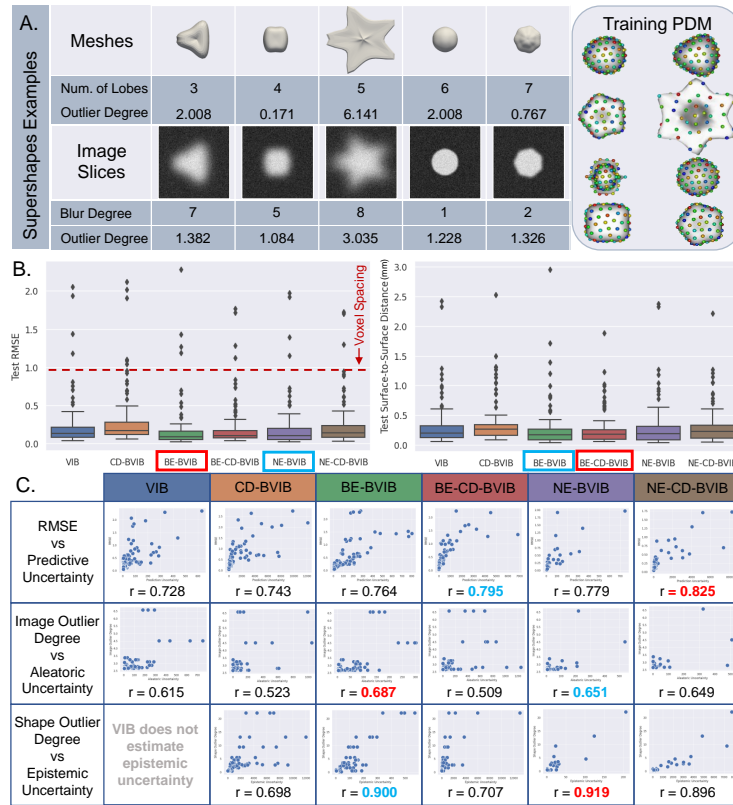


Fig. 2: **Supershapes** (A) Left: Five examples of SS mesh and image pairs with corresponding outlier degrees. Right: Examples of training points overlaid on meshes, where color denotes point correspondence. (B) Box plots show the distribution of errors over the test set, lower is better. (C) Scatterplots show uncertainty correlation, where a higher Pearson r coefficient suggests better-calibrated uncertainty. The best values are marked in red, and the second best in blue.

validation set of size 100, and a testing set of size 100. ShapeWorks [8] was used to optimize PDMs of 128 points on the training set. Target PDMs were then optimized for validation and test sets, keeping the training PDMs fixed so that the test set statistics were not captured by the training PDMs.

Figure 2B demonstrates that all BVIB models performed similarly or better than the baseline VIB in terms of RMSE and surface-to-surface distance, with the BE models performing best. Interestingly, the BE models were more accurate than the NE. This effect could result from the random sign initialization of BE fast weights, which increases members diversity. Adding CD hurt the accuracy slightly, likely because the learning task is made more difficult when layer-wise dropout probabilities are added as variational parameters. However, CD is the cheapest way to add epistemic uncertainty and improve prediction uncertainty calibration. Figure 2C demonstrates prediction uncertainty is well-calibrated for all models (with an error correlation greater than 0.7) and NE-CD-BVIB achieves the best correlation. The aleatoric and epistemic uncertainty correlation was similar across models, with the ensemble-based models performing best.

4.2 Left Atrium Experiments

The left atrium (LA) dataset comprises 1041 anonymized LGE MRIs from unique patients. The images were manually segmented at the University of Utah Division of Cardiovascular Medicine with spatial resolution $0.65 \times 0.65 \times 2.5 \text{ mm}^3$, and the endocardium wall was used to cut off pulmonary veins. The images were cropped around the region of interest, then downsampled by a factor of 0.8 for memory purposes. This dataset contains significant shape variations, including overall size, LA appendage size, and pulmonary veins’ number and length. The input images vary widely in intensity and quality, and LA boundaries are blurred and have low contrast with the surrounding structures. Shapes and image pairs with the largest outlier degrees were held out as outlier test sets, resulting in a *shape outlier* test set of 40 and *image outlier* test set of 78. We randomly split the remaining samples (90%, 10%, 10%) to get a training set of 739, a validation set of 92, and an *inlier* test set of 92. The target PDMs were optimized with ShapeWorks [8] to have 1024 particles.

The accuracy and uncertainty calibration analysis in Figure 3B and 3C show similar results to the supershapes experiment. In both experiments, the proposed combination of dropout and ensembling provided the best-calibrated prediction uncertainty, highlighting the benefit of multimodal Bayesian marginalization. Additionally, the proposed combination gave more accurate predictions on the LA outlier test sets, suggesting improved robustness. BE-CD-BVIB provided the best prediction uncertainty for the LA and the second best (just behind NE-CD-BVIB) for the SS. BE-CD-BVIB is a favorable approach as it does not require training multiple models as NE does and requires relatively low memory addition to the base VIB model. Further qualitative LA results are provided in Appendix F in the form of heat maps of the error and uncertainty on test meshes. Here we can see how the uncertainty correlates locally with the error. As expected, both are highest in the LA appendage and pulmonary veins region,

where LA's and the segmentation process vary the most. It is worth noting a standard normal prior was used for $p(\mathbf{z})$ in all models. Defining a more flexible prior, or potentially learning the prior, could provide better results and will be considered in future work.

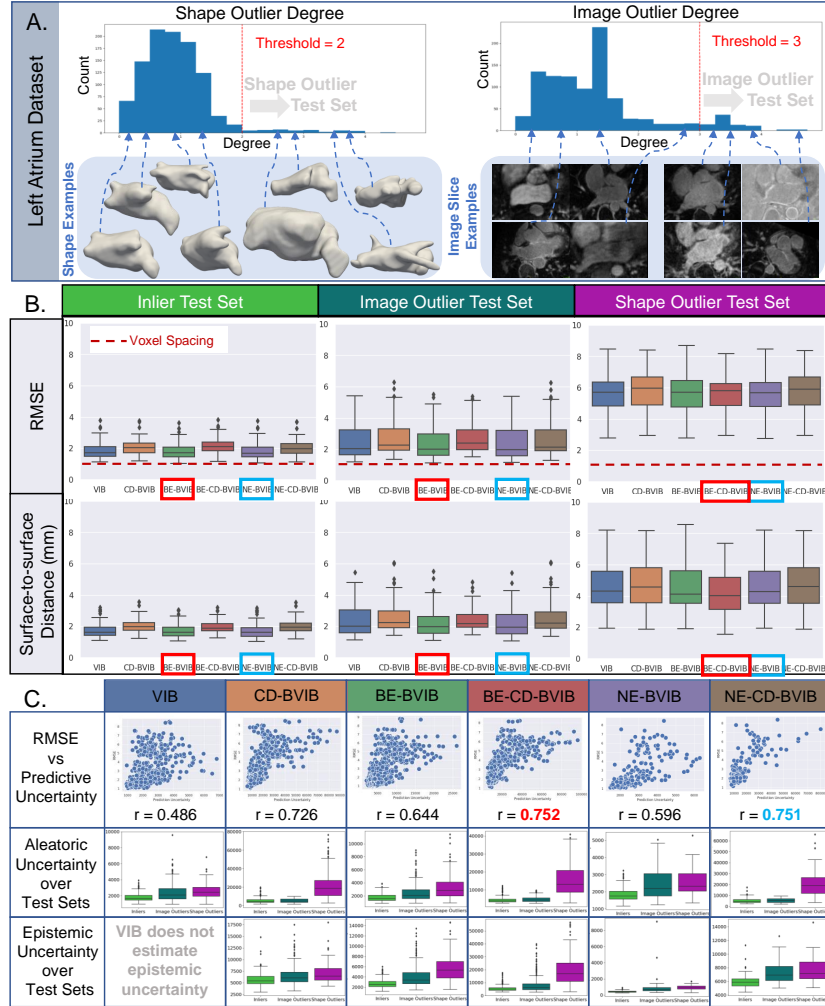


Fig. 3: **Left Atrium** (A) The distribution of shape and image outlier degrees with thresholds is displayed with examples. (B) Box plots show the distribution of errors over the test sets. (C) Scatterplots show uncertainty correlation with error across test sets and box plots show the distribution of uncertainty for each test set. The best values are marked in red, and the second best in blue.

5 Conclusion

The traditional computational pipeline for generating Statistical Shape Models (SSM) is expensive and labor-intensive, which limits its widespread use in clinical research. Deep learning approaches have the potential to overcome these barriers by predicting SSM from unsegmented 3D images in seconds, but such a solution cannot be deployed in a clinical setting without calibrated estimates of epistemic and aleatoric uncertainty. The VIB-DeepSSM model provided a principled approach to quantify aleatoric uncertainty but lacked epistemic uncertainty. To address this limitation, we proposed a fully Bayesian VIB model that can predict anatomical SSM with both forms of uncertainty. We demonstrated the efficacy of two practical and scalable approaches, concrete dropout and batch ensemble, and compared them to the baseline VIB and naive ensembling. Finally, we proposed a novel combination of dropout and ensembling for multimodal Bayesian marginalization and showed that the proposed approach provides improved uncertainty calibration and model robustness on synthetic supershape and real left atrium datasets. These contributions are an important step towards replacing the traditional SSM pipeline with a deep network and increasing the feasibility of fast, accessible SSM in clinical research and practice.

Acknowledgements This work was supported by the National Institutes of Health under grant numbers NIBIB-U24EB029011, NIAMS-R01AR076120, NHLBI-R01HL135568, and NIBIB-R01EB016701. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The authors would like to thank the University of Utah Division of Cardiovascular Medicine for providing left atrium MRI scans and segmentations from the Atrial Fibrillation projects.

References

1. Adams, J., Bhalodia, R., Elhabian, S.: Uncertain-deepssm: From images to probabilistic shape models. *Shape in Medical Imaging : International Workshop, ShapeMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings* **12474**, 57–72 (2020)
2. Adams, J., Elhabian, S.: From images to probabilistic anatomical shapes: A deep variational bottleneck approach. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II*. pp. 474–484. Springer (2022)
3. Alemi, A.A., Fischer, I., Dillon, J.V., Murphy, K.: Deep variational information bottleneck. In: *International Conference on Learning Representations* (2017)
4. Alemi, A.A., Morningstar, W.R., Poole, B., Fischer, I., Dillon, J.V.: Vib is half bayes. In: *Third Symposium on Advances in Approximate Bayesian Inference* (2020)
5. Bhalodia, R., Elhabian, S.Y., Kavan, L., Whitaker, R.T.: Deepssm: A deep learning framework for statistical shape modeling from raw images. In: *Shape In Medical Imaging at MICCAI. Lecture Notes in Computer Science*, vol. 11167, pp. 244–257. Springer (2018)

6. Bhalodia, R., Goparaju, A., Sodergren, T., Whitaker, R.T., Morris, A., Kholmovski, E., Marrouche, N., Cates, J., Elhabian, S.Y.: Deep learning for end-to-end atrial fibrillation recurrence estimation. In: *Computing in Cardiology, CinC 2018*, Maastricht, The Netherlands, September 23-26, 2018 (2018)
7. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: *International conference on machine learning*. pp. 1613–1622. PMLR (2015)
8. Cates, J., Elhabian, S., Whitaker, R.: Shapeworks: Particle-based shape correspondence and visualization software. In: *Statistical Shape and Deformation Analysis*, pp. 257–298. Elsevier (2017)
9. Daxberger, E., Hernández-Lobato, J.M.: Bayesian variational autoencoders for unsupervised out-of-distribution detection. arXiv preprint arXiv:1912.05651 (2019)
10. Der Kiureghian, A., Ditlevsen, O.: Aleatory or epistemic? does it matter? *Structural safety* **31**(2), 105–112 (2009)
11. FDA: Assessing the credibility of computational modeling and simulation in medical device submissions. U.S. Department of Health and Human Services, Food and Drug Administration, Center for Devices and Radiological Health (Dec 2021), <https://www.fda.gov/media/154985/download>
12. Fort, S., Hu, H., Lakshminarayanan, B.: Deep ensembles: A loss landscape perspective. arxiv 2019. arXiv preprint arXiv:1912.02757 (2019)
13. Gal, Y., Hron, J., Kendall, A.: Concrete dropout. *Advances in neural information processing systems* **30** (2017)
14. Gielis, J.: A generic geometric transformation that unifies a wide range of natural and abstract shapes. *American journal of botany* **90**(3), 333–338 (2003)
15. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 9, pp. 249–256. PMLR (13–15 May 2010)
16. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1026–1034 (2015)
17. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems* **30** (2017)
18. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *International Conference on Learning Representations* (12 2014)
19. Maddox, W.J., Izmailov, P., Garipov, T., Vetrov, D.P., Wilson, A.G.: A simple baseline for bayesian uncertainty in deep learning. *Advances in neural information processing systems* **32** (2019)
20. Moghaddam, B., Pentland, A.: Probabilistic visual learning for object representation. *IEEE Transactions on pattern analysis and machine intelligence* **19**(7), 696–710 (1997)
21. Rahaman, R., et al.: Uncertainty quantification and deep ensembles. *Advances in Neural Information Processing Systems* **34**, 20063–20075 (2021)
22. Seitzer, M., Tavakoli, A., Antic, D., Martius, G.: On the pitfalls of heteroscedastic uncertainty estimation with probabilistic neural networks. In: *International Conference on Learning Representations* (2021)
23. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J.J., Shenton, M.E., Gerig, G.: Framework for the statistical shape analysis of brain structures using spharm-pdm. *The insight journal* p. 242 (2006)

24. Thakur, S., Van Hoof, H., Gupta, G., Meger, D.: Unifying variational inference and pac-bayes for supervised learning that scales. arXiv preprint arXiv:1910.10367 (2019)
25. Tóthová, K., Parisot, S., Lee, M.C., Puyol-Antón, E., Koch, L.M., King, A.P., Konukoglu, E., Pollefeys, M.: Uncertainty quantification in cnn-based surface prediction using shape priors. In: Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings. pp. 300–310. Springer (2018)
26. Wen, Y., Tran, D., Ba, J.: Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In: International Conference on Learning Representations (2020)
27. Wilson, A.G., Izmailov, P.: Bayesian deep learning and a probabilistic perspective of generalization. Advances in neural information processing systems **33**, 4697–4708 (2020)
28. Wilson, A.G.: The case for bayesian deep learning. arXiv preprint arXiv:2001.10995 (2020)

A Bayesian VIB Derivation

In Bayesian inference, for a given test sample, \mathbf{x}^* , the predictive distribution is found by marginalizing over Θ : $p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{X}, \mathcal{Y}) = \int p(\mathbf{y}^*|\mathbf{x}^*, \Theta)p(\Theta|\mathcal{X}, \mathcal{Y})d\Theta$. In VIB, for a given test sample \mathbf{x}^* , the predictive distribution is defined as: $p(\mathbf{y}^*|\mathbf{x}^*, \Theta) = \int p(\mathbf{y}^*|\mathbf{z}^*, \theta)p(\mathbf{z}^*|\mathbf{x}^*, \phi)d\mathbf{z}$. To extend VIB to be fully Bayesian, we place a prior ($p(\Theta)$) over model parameters ($\Theta = \{\phi, \theta\}$) and approximate the posterior ($p(\Theta|\mathcal{D})$). Under this framework, we combine the predictive distributions as follows:

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{X}, \mathcal{Y}) = \int \left(\int p(\mathbf{y}^*|\mathbf{z}^*, \theta)p(\mathbf{z}^*|\mathbf{x}^*, \phi)d\mathbf{z} \right) p(\Theta|\mathcal{X}, \mathcal{Y})d\mathbf{z}d\Theta \quad (4)$$

Here we have two intractable posteriors $p(\mathbf{z}|\mathbf{x}, \phi)$ and $p(\Theta|\mathbf{x}, \mathbf{y})$, thus we apply VI twice. The first is approximated via $q(\mathbf{z}|\mathbf{x}, \phi)$ as is done in Equation 2 and the second is approximated by $q(\Theta)$ as is done in Equation 1, resulting in the BVIB objective shown in Equation 3.

B Uncertainty Calculation

To estimate epistemic uncertainty in BVIB-DeepSSM, we generate multiple predictions ($\hat{\mathbf{y}}_n, \hat{\sigma}_n^2$) with T weights sampled from the approximate posterior. In CD, T is the number of different dropout masks used in generating predictions (30). In BE and NE, T is the number of ensemble members (4). In the proposed CD ensemble integration, T is the number of dropout masks times the number of ensemble members (120). The prediction uncertainty is expressed as the sum of the aleatoric and epistemic uncertainty:

$$\Sigma_n^2 = \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{y}}_{n,t}^2 - \left(\frac{1}{T} \sum_{t=1}^T \hat{\mathbf{y}}_{n,t} \right)^2 + \frac{1}{T} \sum_{t=1}^T \hat{\sigma}_{n,t}^2 \quad (5)$$

C Model Parameters

Architecture: The encoder, f_e is comprised of 3D convolution, 3D batch normalization, and fully connected (FC) layers as in [5]. The decoder is comprised of three FC layers and Parametric ReLU activation [16].

Table 1: Parameter values used in all experiments with the tested range.

Parameter	Selected Value	Tested Range/Explanation
β (Equations 2 and 3)	0.01	$[1e^{-8}, 0.9]$
Fixed Learning Rate	$5e^{-5}$	$[1e^{-6}, 1e^{-3}]$
CD Length Scale	$1e^{-3}$	$[1e^{-8}, 1e^2]$
Batch Size	6	Maximum memory allowed.
Number of Ensemble Members	4	Maximum memory allowed.
Adam Optimization Betas	(0.9, 0.999)	Default used.
Adam Optimization Eps	$1e^{-8}$	Default used.
Adam Optimization Weight Decay	0	Default used.

D Example Left Atrium Results

In Figure 4, three examples from the LA test set are provided, including input image slices, ground truth points, and the model output point and uncertainty predictions. Example 1 demonstrates a typical case, example 2 contains an enlarged LA appendage, and example 3 has multiple long pulmonary veins. These examples show the local calibration of the uncertainty estimation, demonstrating the utility of uncertainty quantification in clinical analysis.

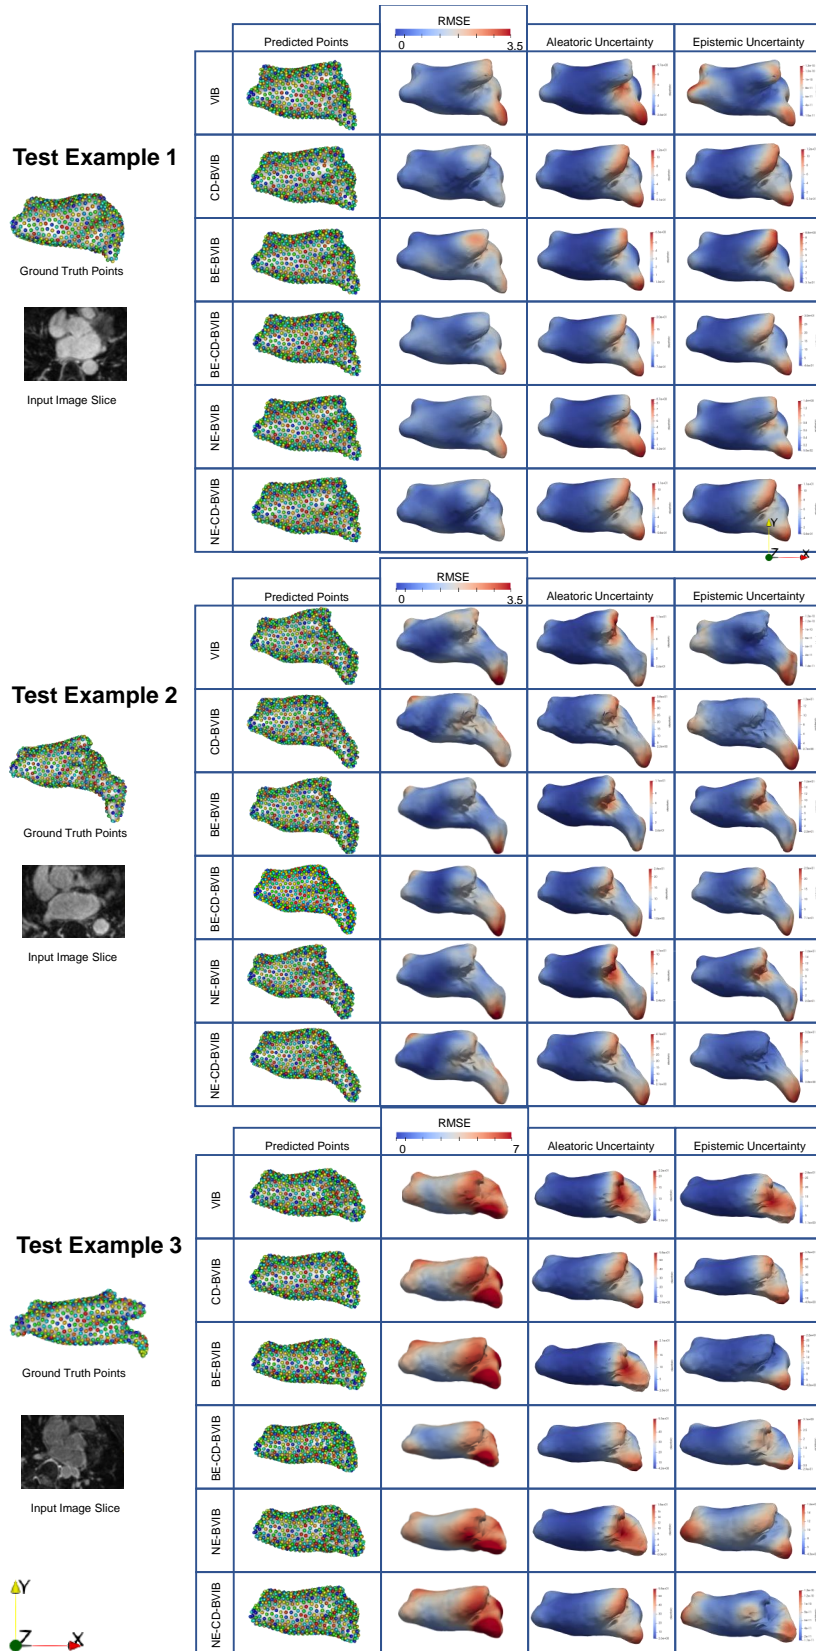


Fig. 4: Three example results from the left atrium test set.