
Deep Deterministic Uncertainty for Semantic Segmentation

Jishnu Mukhoti^{1,2} Joost van Amersfoort¹ Philip H.S. Torr² Yarin Gal¹

Abstract

We extend Deep Deterministic Uncertainty (DDU) (Mukhoti et al., 2021), a method for uncertainty estimation using feature space densities, to semantic segmentation. DDU enables quantifying and disentangling epistemic and aleatoric uncertainty in a single forward pass through the model. We study the similarity of feature representations of pixels at different locations for the same class and conclude that it is feasible to apply DDU location independently, which leads to a significant reduction in memory consumption compared to pixel dependent DDU. Using the DeepLab-v3+ architecture on Pascal VOC 2012, we show that DDU improves upon MC Dropout and Deep Ensembles while being significantly faster to compute.

1. Introduction

With the increasing deployment of deep learning models in safety critical applications like autonomous driving (Huang & Chen, 2020) and medical diagnosis (Esteva et al., 2017), it is imperative for such models to be able to quantify their uncertainty reliably, in addition to making accurate predictions. A significant amount of research has been conducted in this direction and several methods have been introduced in the context of classification (Gal & Ghahramani, 2016; Lakshminarayanan et al., 2017; Blundell et al., 2015). These methods require several forward passes through the model rendering such methods practically infeasible for adoption in large scale applications like semantic segmentation (Long et al., 2015), where dense pixel-wise predictions are necessary, often in real time.

Recently, several methods have been introduced to obtain uncertainty in a single forward pass (Van Amersfoort et al.,

2020; Liu et al., 2020; Mukhoti et al., 2021). In particular, DUQ (Van Amersfoort et al., 2020) and SNGP (Liu et al., 2020) propose using feature extractors with certain inductive biases to impose a bi-Lipschitz constraint on the feature space. They then use a distance aware layer, either an RBF or a Gaussian Process trained end-to-end with the feature extractor. However, both these methods require extensive changes to the model architecture and training setup, with additional hyperparameters which need to be fine-tuned. DDU (Mukhoti et al., 2021) shows that using feature space density with proper inductive biases can capture uncertainty and avoids the problem of *feature collapse* (Van Amersfoort et al., 2020). Due to feature collapse, Out-of-distribution (OoD) samples are often mapped to in-distribution regions in the feature space, making the model overconfident on such inputs. Hence, in order to capture uncertainty through feature space density, one needs to use proper inductive biases on the model architecture.

There are two kinds of uncertainty which are important in deep learning literature: *epistemic uncertainty*, which captures what the model does not know, is high for unseen or OoD inputs and can be reduced with more training data and, *aleatoric uncertainty*, which captures ambiguity and observation noise in in-distribution samples (Kendall & Gal, 2017). In DDU, the epistemic uncertainty is quantified using a feature space density, while the entropy of the softmax distribution can be used to estimate aleatoric uncertainty.

In this paper, we apply and extend DDU to the task of *semantic segmentation* (Long et al., 2015), where each pixel of a given input image is classified to produce an output which has the same spatial dimensions as the input. We choose semantic segmentation in particular as it forms an excellent example of an application with class imbalance and therefore, requires reliable epistemic uncertainty estimates. Furthermore, state-of-the-art models for semantic segmentation (Chen et al., 2017; Zhao et al., 2017; Wang et al., 2020) are large and conventional uncertainty quantification methods like MC Dropout (Gal & Ghahramani, 2016) and Deep Ensembles (Lakshminarayanan et al., 2017) are often prohibitively expensive on such models (Kendall et al., 2015; Mukhoti & Gal, 2018).

The paper is organised as follows: in Section 2, we describe how DDU can be extended for semantic segmentation

^{*}Equal contribution ¹Oxford Applied & Theoretical Machine Learning Group, Department of Computer Science, University of Oxford, Oxford, United Kingdom ²Torr Vision Group, Department of Engineering Science, University of Oxford, Oxford, United Kingdom. Correspondence to: Jishnu Mukhoti <jishnu.mukhoti@eng.ox.ac.uk>.

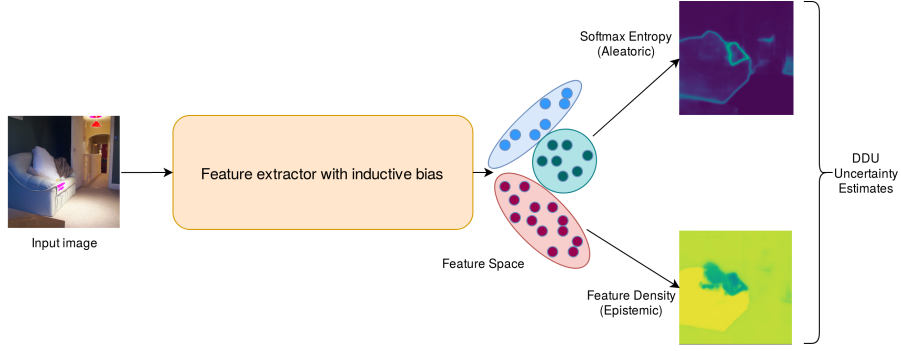


Figure 1. Applying DDU in the context of semantic segmentation

and in Section 3, we provide results using the well-known DeepLab-v3+ (Chen et al., 2017) architecture on the Pascal VOC 2012 (Everingham et al., 2010) dataset to show that DDU outperforms other conventional methods (MC Dropout and Deep Ensembles) of uncertainty quantification in deep learning.

2. DDU in Semantic Segmentation

In this section, we provide details on how DDU can be extended to obtain epistemic and aleatoric uncertainty estimates in semantic segmentation.

A brief introduction to DDU: As described in Mukhoti et al. (2021) in the context of multiclass classification, after training a model with a bi-Lipschitz constraint, we can compute the feature space means and covariances per class using a single pass over all the training samples. The feature space means and covariance matrices can then be used to fit a Gaussian Discriminant Analysis (GDA) (Murphy, 2012). Let z be the feature representation for a given input \mathbf{x} , i.e., $z = f_{\theta}(\mathbf{x})$ where θ represents model parameters. Then the feature density $p(z)$ is computed by marginalizing the density over all classes as

$$p(z) = \sum_c p(z|c)p(c) \quad (1)$$

where $p(z|c)$ is obtained from the GDA and $p(c)$ can be computed directly from the training set. The feature density thus computed can be used to estimate model confidence (opposite of epistemic uncertainty). At the same time, for in-distribution samples, the entropy of the softmax distribution $\mathbb{H}[p(\mathbf{y}|\mathbf{x}, \omega)]$ can be used to capture aleatoric uncertainty.

Pixel-independent class-wise means and covariances: In semantic segmentation, each pixel has a prediction attached to it and a corresponding softmax distribution. A natural question to ask then is whether to compute means and covariance matrices per pixel in order to fit a GDA to the semantic segmentation model. Fortunately, we find that such is not the case and we can indeed compute means




Figure 2. L2 distances between the feature space means of different classes for a pair of distant pixels on the Pascal VOC 2012 val set: (left) Pixels (10, 255) and (500, 255), (middle) Pixels (234, 349) and (36, 22) and (right) Pixels (300, 500) and (400, 255).

and covariance matrices independent of pixels just like in multi-class classification (thereby enforcing invariance). To see this, in Figure 2, we plot the L2 distances between the feature space means of all pairs of classes in the Pascal VOC validation set for two “distant” pixels. We find that the means of the same class are much closer together as compared to other classes irrespective of where the pixels are located. This makes intuitive sense as the convolution kernel (a linear operator) which converts the feature space representations into logits is shared across the entire feature space representation.

Computing feature density: Following the rationale above, we fit a GDA assuming pixels to be independent samples. Hence, we obtain one mean and one covariance matrix per class (not per pixel) and can apply Equation (1) to obtain the feature density per pixel, given an input image. Separately, we can also obtain the per pixel softmax entropy from the model. Using these two we can disentangle aleatoric and epistemic uncertainty with a single deterministic model in semantic segmentation. We present a schematic diagram of this process in Figure 1.

3. Experiments

In this section, we evaluate DDU on semantic segmentation using the well-known Pascal VOC (Everingham et al., 2010) dataset and compare it with three other uncertainty baselines widely applied in practice: softmax entropy, MC Dropout (Gal & Ghahramani, 2016) and Deep Ensembles (Lakshminarayanan et al., 2017).

Baseline	mIoU	Runtime (ms)
Softmax	78.53	275.48 ± 1.91
MC Dropout	78.61	1576.75 ± 1.56
Deep Ensemble	78.47	875.87 ± 0.79
DDU	78.53	263.83 ± 2.79

Table 1. Pascal VOC val set mIoU and runtime in milliseconds of a single forward pass for different baselines averaged over 10 single forward passes. Note that for each single forward pass in the MC Dropout baseline, we perform 5 stochastic forward passes.

Architecture & Training setup: For all the baselines, we use DeepLab-v3+ (Chen et al., 2017) using a ResNet-101 (He et al., 2016) backbone as the architecture of choice. We train each model for 50 epochs on the Pascal VOC training set augmented using the Semantic Boundaries Dataset (SBD) (Hariharan et al., 2011) using SGD as the optimiser with a momentum of 0.9 and a weight decay of $5e-4$. We set the initial learning rate to 0.007 with a polynomial decay during the course of training. Finally, we trained using a batch size of 32 parallelized over 4 GPUs.

Baselines & Uncertainty metrics: As mentioned before, we compare DDU with 3 well-known baselines:

1. **Softmax Entropy**, one of the most commonly used metrics for uncertainty is the entropy $\mathbb{H}[p(\mathbf{y}|\mathbf{x}, \omega)]$ of the softmax distribution $p(\mathbf{y}|\mathbf{x}, \omega)$ (Hendrycks & Gimpel, 2016). This metric is often preferred due to its simplicity and lack of computational overhead. Softmax entropy is known to capture aleatoric uncertainty for in-distribution samples (Mukhoti et al., 2021). However, it cannot capture epistemic uncertainty reliably (eg. for OoD inputs).
2. **MC Dropout (MCDO)** (Gal & Ghahramani, 2016) is a method which uses dropout at test time as an approximation to Bayesian inference. Multiple stochastic forward passes are performed with dropout layers active during test time. The softmax distributions $p(\mathbf{y}|\mathbf{x}, \hat{\omega})$ obtained from these forward passes can then be used to compute either *predictive entropy* (PE): $\mathbb{H}[\mathbb{E}_{\hat{\omega}}[p(\mathbf{y}|\mathbf{x}, \hat{\omega})]]$ or *mutual information* (MI): $\mathbb{H}[\mathbb{E}_{\hat{\omega}}[p(\mathbf{y}|\mathbf{x}, \hat{\omega})]] - \mathbb{E}_{\hat{\omega}}[\mathbb{H}[p(\mathbf{y}|\mathbf{x}, \hat{\omega})]]$ (Houlsby et al., 2011) as measures of uncertainty. While MI is known to estimate epistemic uncertainty, PE captures both epistemic and aleatoric uncertainty (Gal, 2016). In our experiments, we implement MC dropout by activating the dropout layers in the DeepLab-v3+ architecture during test time. We don’t insert new dropout layers. Finally, we use 5 stochastic forward passes for MC Dropout.
3. **Deep Ensembles** (Lakshminarayanan et al., 2017) is a simple method where an ensemble of neural networks

is trained. Similar to MC Dropout, both the PE as well as the MI from the ensemble predictions can be used to estimate uncertainty. In our experiments we use an ensemble of 3 DeepLab-v3+ models, all trained with identical architecture and training setup.

Metrics for evaluation: In order to evaluate the quality of uncertainty in semantic segmentation, we use the metrics proposed in (Mukhoti & Gal, 2018): $p(\text{accurate}—\text{certain})$, $p(\text{uncertain}—\text{inaccurate})$ and PAVPU. The metric $p(\text{accurate}—\text{certain})$ measures the probability of a prediction being accurate given that the model is confident on the prediction. Similarly, $p(\text{uncertain}—\text{inaccurate})$ measures the probability of the model being uncertain on inaccurate predictions. PAVPU computes the probability of the model being confident on an accurate prediction or uncertain on an inaccurate one. A good model should ideally have high values on all these 3 metrics. Note that these metrics depend on a threshold for uncertainty, i.e., to define a prediction as certain or uncertain. Hence, they can be computed for different uncertainty thresholds. We plot the performance of all the baselines on these metrics for different uncertainty thresholds in Figure 3.

In addition, we report the Pascal VOC 2012 validation set accuracy and the runtime of a single forward pass for all the baselines in Table 1. Note that a single forward pass for the MC Dropout baseline consists of five stochastic forward passes and a single forward pass from the ensemble involves getting predictions from three ensemble components. Finally, we visualise the uncertainty estimates from each baseline for four samples from the Pascal VOC 2012 val set in Figure 4.

Observations: Firstly we note from Table 1 that the runtime of DDU and a normal softmax model with values around 263 and 275 milliseconds respectively, are far lower than MC Dropout and Deep Ensembles. In fact, a single forward pass in MC Dropout requires around 1.57 seconds on an Nvidia Quadro RTX 6000 GPU. Although these baselines have not been tuned for runtime, real-time latency requirements of around 200ms (i.e., 5 predictions a second) make adoption of time-consuming methods infeasible in real-life applications. Furthermore, note that the val set mIoU for all the models are very similar.

Secondly, from Figure 3, we can see that DDU, having higher values on $p(\text{accurate}—\text{certain})$, $p(\text{uncertain}—\text{inaccurate})$ and PAVPU for most uncertainty thresholds, outperforms all other baselines on all these three metrics.

Finally, from Figure 4, we note that DDU feature density captures epistemic uncertainty whereas softmax entropy captures aleatoric uncertainty. Note that aleatoric uncertainty is high on edges of objects as those are the regions with




Figure 3. Evaluation metrics: $p(\text{accurate}|\text{certain})$, $p(\text{uncertain}|\text{inaccurate})$ and PAVPU evaluated on different baselines on the PASCAL VOC validation set. DDU outperforms all other baselines.

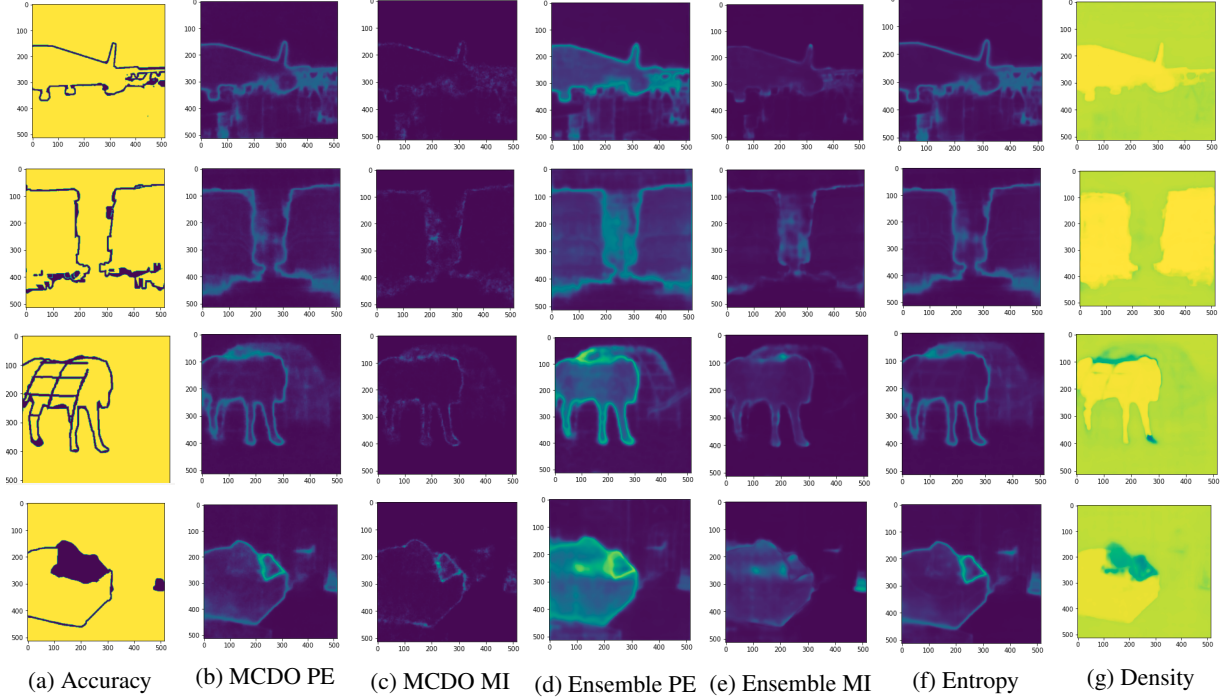


Figure 4. Visualisation of different uncertainty baselines on samples from the PASCAL VOC validation set. The first column captures pixel-wise accuracy with bright signifying accurate and dark, inaccurate. The second and third columns show predictive entropy (PE) and mutual information (MI) obtained from the MC Dropout (MCDO) baseline respectively, the fourth and fifth columns show the PE and MI from deep ensembles. The sixth column maps per-pixel softmax entropy which is the aleatoric uncertainty estimate of DDU, and finally the seventh column shows feature density, which is the epistemic component captured by DDU. All the baselines save DDU density (last column on the right) capture uncertainty, i.e., the brighter, the more uncertain whereas DDU feature density captures confidence and hence brighter pixels signify more confident pixels and vice versa.

maximum ambiguity or observation noise. On the other hand, epistemic uncertainty is high on regions which are previously unseen (or less seen) by the model. In the first two samples (first two rows), the epistemic uncertainty is not high and aleatoric uncertainty is captured on the edges by softmax entropy. In the last sample (last row), the epistemic uncertainty is high for a big patch which is inaccurately predicted. DDU feature density for that entire patch is significantly lower whereas softmax entropy doesn't capture it and is only high on the edges.

4. Conclusion

In this paper, we show that Deep Deterministic Uncertainty (DDU) can be easily extended to the task of semantic segmentation. We find that fitting DDU to a semantic segmentation model with a fully convolutional architecture can be done in a pixel-independent fashion, thereby making its adoption relatively simple. Finally, with experiments on Pascal VOC 2012 using DeepLab-v3+, we observe that DDU outperforms other well-known methods of uncertainty quantification without compromising on accuracy/mIoU and with the runtime of a single deterministic model.

References

- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622. PMLR, 2015.
- Chen, L.-C., Papandreou, G., Schroff, F., and Adam, H. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., and Thrun, S. Dermatologist-level classification of skin cancer with deep neural networks. *nature*, 542(7639):115–118, 2017.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88 (2):303–338, 2010.
- Gal, Y. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
- Hariharan, B., Arbelaez, P., Bourdev, L., Maji, S., and Malik, J. Semantic contours from inverse detectors. In *International Conference on Computer Vision (ICCV)*, 2011.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Houlsby, N., Huszár, F., Ghahramani, Z., and Lengyel, M. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- Huang, Y. and Chen, Y. Autonomous driving with deep learning: a survey of state-of-art technologies. *arXiv preprint arXiv:2006.06091*, 2020.
- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? *arXiv preprint arXiv:1703.04977*, 2017.
- Kendall, A., Badrinarayanan, V., and Cipolla, R. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv preprint arXiv:1511.02680*, 2015.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf>.
- Liu, J. Z., Lin, Z., Padhy, S., Tran, D., Bedrax-Weiss, T., and Lakshminarayanan, B. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *arXiv preprint arXiv:2006.10108*, 2020.
- Long, J., Shelhamer, E., and Darrell, T. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Mukhoti, J. and Gal, Y. Evaluating bayesian deep learning methods for semantic segmentation. *arXiv preprint arXiv:1811.12709*, 2018.
- Mukhoti, J., Kirsch, A., van Amersfoort, J., Torr, P. H., and Gal, Y. Deterministic neural networks with appropriate inductive biases capture epistemic and aleatoric uncertainty. *arXiv preprint arXiv:2102.11582*, 2021.
- Murphy, K. P. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Van Amersfoort, J., Smith, L., Teh, Y. W., and Gal, Y. Uncertainty estimation using a single deep deterministic neural network. In *International Conference on Machine Learning*, pp. 9690–9700. PMLR, 2020.
- Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- Zhao, H., Shi, J., Qi, X., Wang, X., and Jia, J. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2881–2890, 2017.