
Depth Uncertainty Networks for Active Learning

Chelsea Murray James U. Allingham Javier Antorán José Miguel Hernández-Lobato
Department of Engineering
University of Cambridge
{c1m88, jua23, ja666, jmh233}@cam.ac.uk

Abstract

In active learning, the size and complexity of the training dataset changes over time. Simple models that are well specified by the amount of data available at the start of active learning might suffer from bias as more points are actively sampled. Flexible models that might be well suited to the full dataset can suffer from overfitting towards the start of active learning. We tackle this problem using Depth Uncertainty Networks (DUNs), a BNN variant in which the depth of the network, and thus its complexity, is inferred. We find that DUNs outperform other BNN variants on several active learning tasks. Importantly, we show that on the tasks in which DUNs perform best they present notably less overfitting than baselines.

1 Introduction

In many important applications of machine learning, labelled data can be scarce and expensive to obtain. Active learning is one approach to improving data efficiency in such cases. Given a fixed labelling budget, the objective of active learning is to identify examples that maximise the expected gain in model performance. Bayesian Neural Networks (BNNs) are a natural fit for active learning problems as they provide uncertainty estimates that can be used as a measure of the informativeness of unlabelled data points. This work focuses on a recently proposed form of BNN, Depth Uncertainty Networks (DUNs), in which probabilistic inference is performed over the depth of the network [Antorán et al., 2020]. We hypothesise that the ability to infer depth enables DUNs to adapt model complexity to the varying size of the training dataset as additional labels are acquired. This is expected to result in reduced overfitting in DUNs in the early stages of active learning, and in better performance of DUNs relative to other BNN methods.

2 Background

2.1 Active learning

In the active learning framework, a model is initially trained on a small labelled subset of the available data, and additional unlabelled points are selected via an *acquisition function* to be labelled by an external *oracle* (e.g. a human expert) [Settles, 2010]. Given a model with parameters θ trained on training data $\mathcal{D}_{\text{train}} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^N$, the acquisition function $\alpha(\cdot)$ scores all unlabelled examples in the pool set $\mathcal{D}_{\text{pool}} = \{\mathbf{x}_j\}_{j=1}^{N_{\text{pool}}}$. These scores are used to select the next point \mathbf{x}^* to be labelled:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{D}_{\text{pool}}} \alpha(\mathbf{x}; \theta, \mathcal{D}_{\text{train}}). \quad (1)$$

Houlsby et al. [2011] propose an acquisition called Bayesian Active Learning by Disagreement (BALD),

$$\alpha_{\text{BALD}}(\mathbf{x}; \theta, \mathcal{D}_{\text{train}}) = \mathbb{H}[\mathbf{y} | \mathbf{x}, \mathcal{D}_{\text{train}}] - \mathbb{E}_{\theta \sim p(\theta | \mathcal{D}_{\text{train}})}[\mathbb{H}[\mathbf{y} | \mathbf{x}, \theta]], \quad (2)$$

which selects points for which the predictions of individual parameterisations maximally disagree—i.e., where there is high uncertainty in the predictive posterior on average—but the predictions of individual parameter settings are confident. To avoid acquiring correlated points when performing batch acquisition, we implement a stochastic relaxation of BALD,

$$\mathbf{x}^* \sim \alpha_{\text{BALDstoch}}(\mathbf{x}; \theta, \mathcal{D}_{\text{train}}) = \text{softmax}(T\alpha_{\text{BALD}}(\mathbf{x}; \theta, \mathcal{D}_{\text{train}})). \quad (3)$$

We choose temperature $T = 10$; see appendix D.4 for a discussion of this choice.

2.2 Depth uncertainty networks

DUNs, depicted in figure 1, are a BNN variant in which the depth of the network $d \in [1, 2, \dots, D]$ is treated as a random variable. Model weights θ are kept deterministic, simplifying optimisation. We place a categorical prior over depth $p(d) = \text{Cat}(d|\{\beta_i\}_{i=0}^D)$, resulting in a categorical posterior which can be marginalised in closed form. This procedure can be seen as Bayesian model averaging over an ensemble of subnetworks of increasing depth. Thus, for DUNs, the predictive distribution and $\alpha_{\text{BALD}}(\cdot)$ are tractable and cheap to compute. See appendix A or [Antorán et al., 2020] for a detailed description.

In DUNs, predictions from different depths of the network correspond to different kinds of functions—shallow networks induce simple functions, while deeper networks induce more complex functions. We hypothesise that DUNs will automatically adapt model flexibility during active learning by inferring initially shallow networks, and progressively deeper networks as more data are observed. This capability is expected to reduce overfitting in the small data regimes typical of active learning, and result in improved performance for DUNs relative to other BNN methods. It is worth noting that automatically adapting model complexity to observed data is also a promised feature of traditional BNNs. However, unlike in DUNs where inference is exact, weight space models require crude approximations which fall short in this regard [Ghosh et al., 2019].




Figure 1: DUN graphical model.

3 Results

We perform experiments on five toy regression datasets, nine UCI regression datasets [Hernández-Lobato and Adams, 2015], and MNIST. Concrete, Energy, and Kin8nm are presented in the following sections, with the rest provided in appendix D. The experimental setup is described in appendix C. Experiments are repeated 40 times, with the mean and standard deviations reported in our figures.

3.1 DUN posteriors

First, we investigate whether our hypothesis that DUNs adapt their inferred depth to the size of the dataset holds. Figure 2 compares the posterior probabilities over depth for DUNs trained on datasets from the first and final steps of active learning, and illustrates that this does occur in practice.

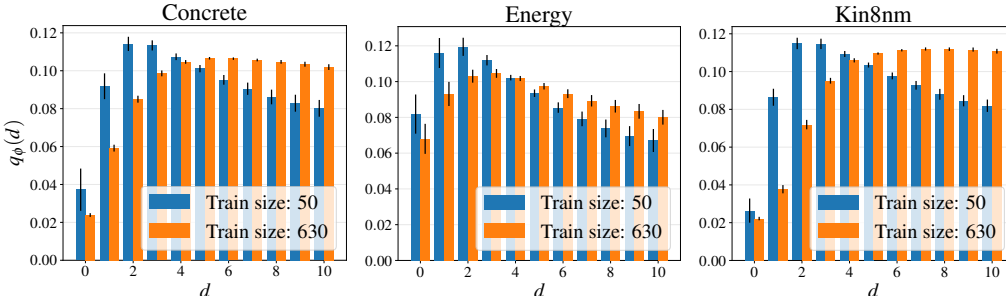


Figure 2: DUN depth posteriors for three UCI regression datasets, with the **smallest** and **largest** labelled datasets used during active learning. Additional datasets’ results are in figures D.1 and D.4c.

3.2 DUNs versus other BNN methods

We compare the negative log-likelihood (NLL) performance of DUNs to mean-field variational inference (MFVI) [Blundell et al., 2015, Graves, 2011] and Monte Carlo dropout (MCDO) [Gal and

Ghahramani, 2016] in figure 3. The test NLL for MFVI and MCDO is evaluated using 10 Monte Carlo samples. DUNs perform either better than, or comparably to, MCDO, and clearly outperform MFVI. Following an approach proposed by Farquhar et al. [2021] for quantifying overfitting in the context of active learning problems, figure 4 shows that the magnitude of overfitting in DUNs tends to be smaller relative to MCDO. Given the results in figure 2 and figure 4, it is plausible to conclude that DUNs’ flexibility over depth is useful in active learning settings. DUNs’ performance may also be attributable to better quality uncertainty estimates, as reported by Antorán et al. [2020], or to the fact that α_{BALD} can be computed exactly for DUNs.




Figure 3: Test NLL vs. number of training points using stochastic BALD acquisition function evaluated on UCI datasets. **DUNs**, **MCDO** and **MFVI** are compared. Additional datasets’ results are in figures D.2, D.4b and D.5. Results for classification error rather than NLL are in figure D.6.




Figure 4: Evolution of “overfitting bias”, as described by Farquhar et al. [2021]. See appendix B for more details. For **DUNs** and **MCDO** as the amount of acquired data M grows. DUNs are more robust to overfitting than MCDO. Results for additional datasets are in figure D.3.

4 Discussion

Our preliminary investigation into the application of DUNs to active learning has led to some promising results. DUNs do, as hypothesised, change their complexity as the dataset grows. As a result of this property, DUNs not only exhibit improved performance—on 10 of the 15 datasets analysed—but also overfit to a lesser extent relative to alternative BNN methods. Overfitting tends to be lower for DUNs only for the datasets on which DUNs also outperform MCDO, suggesting that robustness to overfitting is an important factor in DUNs’ superior performance. Interestingly, DUNs tend to outperform MCDO more as the number of actively sampled points increases. Intuitively, reduced overfitting should have the greatest impact for the smallest training set sizes. This may be explained by DUNs having better quality uncertainty estimates [Antorán et al., 2020], which are more important later in active learning when it is more difficult to differentiate between the informativeness of candidate data points. Alternatively, the extra depth of the DUNs may enable better performance as the dataset becomes more complex. Or, perhaps the negative effects of overfitting early in the active selection process accumulate and have a bigger influence later on.

We suggest two avenues for further investigation. Firstly, the observations made in this work should be confirmed using more complex datasets (preliminary experiments were also conducted on MNIST, with results shown in figure D.4). We also conducted investigations into the effect of using decaying prior probabilities over depth in place of a uniform prior, but found that this has almost no impact on performance (figure D.8). Further work could investigate the development of an improved DUN in which the prior over depth plays a larger role in controlling model complexity.

References

- Javier Antorán, James Allingham, and José Miguel Hernández-Lobato. Depth uncertainty in neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- Javier Antorán, James Urquhart Allingham, and José Miguel Hernández-Lobato. Variational depth search in ResNets. *CoRR*, abs/2002.02797, 2020.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- A D’Amour, K Heller, D Moldovan, B Adlam, B Alipanahi, A Beutel, C Chen, J Deaton, J Eisenstein, MD Hoffman, et al. Underspecification presents challenges for credibility in modern machine learning. arxiv 2020. *arXiv preprint arXiv:2011.03395*, 2020.
- Sebastian Farquhar, Yarin Gal, and Tom Rainforth. On statistical bias in active learning: How and when to fix it. *International Conference on Learning Representations*, 2021.
- Andrew Y. K. Foong, Yingzhen Li, José Miguel Hernández-Lobato, and Richard E. Turner. ‘In-between’ uncertainty in Bayesian neural networks. *CoRR*, abs/1906.11537, 2019.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016.
- Soumya Ghosh, Jiayu Yao, and Finale Doshi-Velez. Model selection in bayesian neural networks via horseshoe priors. *Journal of Machine Learning Research*, 20(182):1–46, 2019. URL <http://jmlr.org/papers/v20/19-236.html>.
- Alex Graves. Practical variational inference for neural networks. In *Advances in neural information processing systems*, pages 2348–2356. Citeseer, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *International Conference on Machine Learning*, pages 1861–1869. PMLR, 2015.
- Neil Houlsby, Ferenc Huszar, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *CoRR*, abs/1112.5745, 2011. URL <http://arxiv.org/abs/1112.5745>.
- S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Pavel Izmailov, Wesley J Maddox, Polina Kirichenko, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Subspace inference for Bayesian deep learning. In *Uncertainty in Artificial Intelligence*, pages 1169–1179. PMLR, 2020.
- Durk P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/bc7316929fe1545bf0b98d114ee3ecb8-Paper.pdf>.
- Yann LeCun, Corinna Cortes, and C Burges. MNIST handwritten digit database, 1998. URL <http://www.research.att.com/~yann/ocr/mnist>, 1998.
- Burr Settles. Active learning literature survey. *Machine Learning*, 15(2):201–221, 2010.
- Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I*, volume 8689 of *Lecture Notes in Computer Science*, pages 818–833. Springer, 2014.

A Depth uncertainty networks

BNNs translate uncertainty in weight-space into predictive uncertainty by marginalising out the posterior over weights. The intractability of this posterior necessitates the use of approximate inference methods. DUNs, in contrast, leverage uncertainty about the appropriate *depth* of the network in order to obtain predictive uncertainty estimates. There are two primary advantages of this approach: 1) the posterior over depth is tractable, mitigating the need for limiting approximations; and 2) due to the sequential structure of feed-forward neural networks, both inference and prediction can be performed with a single forward pass, making DUNs suitable for deployment in low-resource environments [Antorán et al., 2020].

DUNs are composed of subnetworks of increasing depth, with each subnetwork contributing one prediction to the final model. The computational model for DUNs is shown in figure A.1. Inputs are passed through an input block, $f_0(\cdot)$, and then sequentially through each of D intermediate blocks $\{f_i(\cdot)\}_{i=1}^D$, with the activations of the previous layer acting as the inputs of the current layer. The activations of each of the $D + 1$ layers are passed through the output block $f_{D+1}(\cdot)$ to generate per-depth predictions. These are combined via marginalisation over the depth posterior, equivalent to forming an ensemble with networks of increasing depth. Variation between the predictions from each depth can be used to obtain predictive uncertainty estimates, in much the same way that variance in the predictions of different sampled weight configurations yield predictive uncertainties in BNNs.




Figure A.1: DUN computational model [Antorán et al., 2020]. Each layer’s activations are passed through the output block, producing per-depth predictions.

A.1 Uncertainty over depth

Figure A.2 compares the graphical model representations of a BNN and a DUN. In BNNs, the weights θ are treated as random variables drawn from a distribution $p_\gamma(\theta)$ with hyperparameters γ . In DUNs, the depth of the network d is assumed to be stochastic, while the weights are learned as hyperparameters. A categorical prior distribution is assumed for d , with hyperparameters β : $p_\beta(d) = \text{Cat}(d | \{\beta_i\}_{i=0}^D)$. The model’s marginal log likelihood (MLL) is given by

$$\log p(\mathcal{D}_{\text{train}}; \theta) = \log \sum_{i=0}^D \left(p_\beta(d = i) \prod_{n=1}^N p(\mathbf{y}_n | \mathbf{x}_n, d = i; \theta) \right), \quad (4)$$

where the likelihood for each depth is parameterised by the corresponding subnetwork’s output: $p(\mathbf{y} | \mathbf{x}, d = i; \theta) = p(\mathbf{y} | f_{D+1}(\mathbf{a}_i; \theta))$, where $\mathbf{a}_i = f_i(\mathbf{a}_{i-1})$.



Figure A.2: Left: graphical model of a BNN. Right: graphical model of a DUN.

The posterior,

$$p(d \mid \mathcal{D}_{\text{train}}; \theta) = \frac{p(\mathcal{D}_{\text{train}} \mid d; \theta)p_{\beta}(d)}{p(\mathcal{D}_{\text{train}}; \theta)}, \quad (5)$$

is also a categorical distribution, whose probabilities reflect how well each depth subnetwork explains the data. DUNs leverage two properties of modern neural networks: first, that successive layers have been shown to extract features at increasing levels of abstraction [Zeiler and Fergus, 2014]; and second, that networks are typically underspecified, meaning that several different models may explain the data well [D’Amour et al., 2020]. The first property implies that initial layers in a network specialise on low-level feature extraction, yielding poor predictions from the shallower subnetworks in a DUN. This is handled by the depth posterior assigning low probabilities to earlier layers, in preference for later layers that specialise on prediction. The second property suggests that subnetworks at several depths are able to explain the data simultaneously and thus have non-zero posterior probabilities, yielding the diversity in predictions required to obtain useful estimates of model uncertainty.

A.2 Inference in DUNs

Antorán et al. [2020] find that directly optimising the MLL equation (4) with respect to θ results in convergence to a local optimum in which all but one layer is attributed a posterior probability of zero. That is, direct optimisation returns a deterministic network of arbitrary depth. The authors instead propose a stochastic gradient variational inference (VI) approach with the aim of separating optimisation of the weights θ from the posterior. A surrogate categorical distribution $q_{\phi}(d)$ is introduced as the variational posterior. The following evidence lower bound (ELBO) can be derived:

$$\mathcal{L}(\phi, \theta) = \sum_{n=1}^N \mathbb{E}_{q_{\phi}(d)} [\log p(\mathbf{y}_n \mid \mathbf{x}_n, d; \theta)] - \text{KL}(q_{\phi}(d) \parallel p_{\beta}(d)), \quad (6)$$

allowing θ and the variational parameters ϕ to be optimised simultaneously. Antorán et al. [2020] show that under the VI approach, the variational parameters converge more slowly than the weights, enabling the weights to reach a setting at which a positive posterior probability is learnt for several depths.

It is important to note that equation (6) can be computed exactly, and that optimising $\mathcal{L}(\phi, \theta)$ is equivalent to exact inference of the true posterior $p(d \mid \mathcal{D}_{\text{train}}; \theta)$ in the limit: since both q and p are categorical, equation (6) is convex and tight at the optimum (i.e., $q_{\phi}(d) = p(d \mid \mathcal{D}_{\text{train}}; \theta)$). Since the expectation term can be computed exactly using the activations at each layer (recall $p(\mathbf{y} \mid \mathbf{x}, d = i; \theta) = p(\mathbf{y} \mid f_{D+1}(\mathbf{a}_i; \theta))$), the ELBO can be evaluated exactly without Monte Carlo sampling.

B Bias in active learning

We use notation similar to Farquhar et al. [2021]. In the active learning setting, given a loss function $\mathcal{L}(\mathbf{y}, f_\theta(\mathbf{x}))$, we aim to find θ that minimises the population risk over $p_{\text{data}}(\mathbf{y}, \mathbf{x})$: $\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_{\text{data}}} [\mathcal{L}(\mathbf{y}, f_\theta(\mathbf{x}))]$. In practice, we only have access to N samples from p_{data} . These yield the empirical risk, $r = \frac{1}{N} \sum_{n=1}^N \mathcal{L}(\mathbf{y}_n, f_\theta(\mathbf{x}_n))$. r is an unbiased estimator of the population risk when the data are drawn i.i.d. from p_{data} . If these samples have been used to train θ , our estimators become biased. We refer to estimators biased by computing with the training data with capital letters (i.e., R). In the active learning setting, our model is optimised using a subset of M actively sampled data points:

$$\tilde{R} = \frac{1}{M} \sum_{m=1}^M \mathcal{L}(\mathbf{y}_{i_m}, f_\theta(\mathbf{x}_{i_m})), \quad i_m \sim \alpha(i_{1:m-1}, \mathcal{D}_{\text{pool}}). \quad (7)$$

The proposal distribution $\alpha(i_m; i_{1:m-1}, \mathcal{D}_{\text{pool}})$ over the pool of unlabelled data $\mathcal{D}_{\text{pool}}$ represents the probability of each index being sampled next, given that we have already acquired $\mathcal{D}_{\text{train}} = \{\mathbf{x}_i\}_{1:m-1}$. As a result of active sampling, \tilde{R} is not an unbiased estimator of R . Farquhar et al. [2021] propose \tilde{R}_{LURE} (“levelled unbiased risk estimator”), which removes active learning bias:

$$\tilde{R}_{\text{LURE}} \equiv \frac{1}{M} \sum_{m=1}^M v_m \mathcal{L}_{i_m}; \quad v_m \equiv 1 + \frac{N-M}{N-m} \left(\frac{1}{(N-m+1) \alpha(i_m; i_{1:m-1}, \mathcal{D}_{\text{pool}})} - 1 \right), \quad (8)$$

where $\mathcal{L}_{i_m} \equiv \mathcal{L}(\mathbf{y}_{i_m}, f_\theta(\mathbf{x}_{i_m}))$. Intuitively, the estimator works by re-weighting each example’s contribution to the total loss by its inverse acquisition probability, such that if unusual examples have particularly high acquisition probabilities, they contribute proportionally less to the total risk. A detailed derivation of this estimator is given by Farquhar et al. [2021].

Surprisingly, Farquhar et al. [2021] find that using \tilde{R}_{LURE} during training can negatively impact the performance of flexible models. They explain this by noting that the overfitting bias ($r-R$) typical in such models has the opposite sign to the bias induced by active learning. That is, active learning encourages the selection of “difficult” to predict data points, increasing the overall risk, whereas overfitting results in the total risk being underestimated. Where overfitting is substantial, it is (partially) offset by active learning bias, such that correcting for active learning bias does not improve model performance.

In this paper, however, we do not attempt to use \tilde{R}_{LURE} to improve the performance of active learning, but rather as a means to measure the degree of overfitting for DUNs and MCDO. The procedure is as follows:

1. We train a model on the current active learning train set.
2. We calculate the test-set empirical risk r of the model.
3. We calculate the unbiased current train-set empirical risk \tilde{R}_{LURE} of the model. This quantity is free from active learning bias.
4. We estimate the degree of overfitting as the difference between these two values of the empirical risk. Intuitively, once active learning bias has been accounted for, the difference between the risk estimated on the train and test sets must be due to overfitting.

C Experimental setup

We conduct experiments on three groups of datasets: the toy regression datasets used in Antorán et al. [2020], nine of the UCI regression datasets [Hernández-Lobato and Adams, 2015], and the MNIST image classification dataset [LeCun et al., 1998]. We implement batch-based active learning, with batches of 10 (for MNIST and most toy datasets) or 20 (for most UCI datasets) data points acquired in each query, depending on the size of the dataset. An initial training set representing a small proportion of the full data is selected uniformly at random in the first query. For all regression datasets, 80% of the data are used for training, 10% for validation and 10% for testing. The standard train-test split is used for MNIST. Details about dataset sizes, input dimensionality, and active learning specifications for each dataset are provided in table C.1.

Table C.1: Summary of datasets and active learning specifications. For the toy and UCI datasets, 80% of the data are used for training, 10% for validation and 10% for testing. For MNIST, the test and train set sizes are shown in parentheses, i.e., (train & test).

NAME	SIZE	INPUT DIM.	INIT. TRAIN SIZE	NO. QUERIES	QUERY SIZE
Simple 1d	501	1	10	30	10
Izmailov et al. [2020]	400	1	10	30	10
Foong et al. [2019]	100	1	10	15	5
Matern	400	1	10	30	10
Wiggle	300	1	10	20	10
Boston Housing	506	13	20	17	20
Concrete Strength	1,030	8	50	30	20
Energy Efficiency	768	8	50	30	20
Kin8nm	8,192	8	50	30	20
Naval Propulsion	11,934	16	50	30	20
Power Plant	9,568	4	50	30	20
Protein Structure	45,730	9	50	30	20
Wine Quality Red	1,599	11	50	30	20
Yacht Hydrodynamics	308	6	20	20	10
MNIST	70,000 (60,000 & 10,000)	784 (28 × 28)	10	10	10

For the regression problems, we implement a fully-connected network with residual connections, with 100 hidden nodes per layer. The networks contain 10 hidden layers for DUNs, or three hidden layers for MCDO and MFVI. Other depths were tested for the baseline methods, with similar results to the chosen depth of three layers (see figure D.9 for these results). We use ReLU activations, and for DUNs batch normalisation is applied after every layer [Ioffe and Szegedy, 2015]. Optimisation is performed over 1,000 iterations with full-batch gradient descent, with momentum of 0.9 and a learning rate of 10^{-4} . A weight decay value of 10^{-5} is also used. We do not implement early stopping, but the best model based on evaluation of the evidence lower bound on the validation set is used for reporting final results. MFVI models are trained using five MC samples and the local reparameterisation trick [Kingma et al., 2015], and prediction for both MCDO and MFVI is based on 10 MC samples. For MCDO models a fixed dropout probability of 0.1 is used. Unless otherwise specified, DUNs use a uniform categorical prior over depth, while MFVI networks use a $\mathcal{N}(\mathbf{0}, I)$ prior over weights. These hyperparameter settings largely follow those used in Antorán et al. [2020] for the toy regression problems.

For image classification experiments we use the convolutional network architecture described in Antorán et al. [2020]. DUNs contain 10 intermediate blocks (bottleneck convolutional blocks described in He et al. [2016]), while MCDO and MFVI networks contain three. Following PyTorch defaults for MNIST, optimisation is performed over 90 epochs, with the learning rate reduced by a factor of 10 at epochs 40 and 70. The initial learning rate is 0.1. We do not use a validation set to find the best model, instead reporting results on the final model. All other hyperparameters are as for the regression case.

All experiments are repeated 40 times with different weight initialisations and train-test splits (with the exception of experiments on the Protein dataset, which are repeated only 30 times due to the cost of evaluation on the larger test set). Unless otherwise specified, we report the mean and standard deviation of the relevant metric over the repeated experiment runs.

D Additional results

D.1 Remaining UCI regression datasets

This section presents extensions of the results shown in the main text of this work, for the remaining six of the nine UCI regression datasets investigated. Figure D.1 shows the DUN posteriors, as in figure 2. Figure D.2 compares the performance of DUNs on active learning problems to MCDO and MFVI, as in figure 3. Figure D.3 compares overfitting in DUNs and MCDO, as in figure 4.




Figure D.1: Posterior probabilities over depth for DUNs trained on UCI regression datasets, for the **smallest** and **largest** labelled datasets used in active learning. Extension of the results in figure 2.

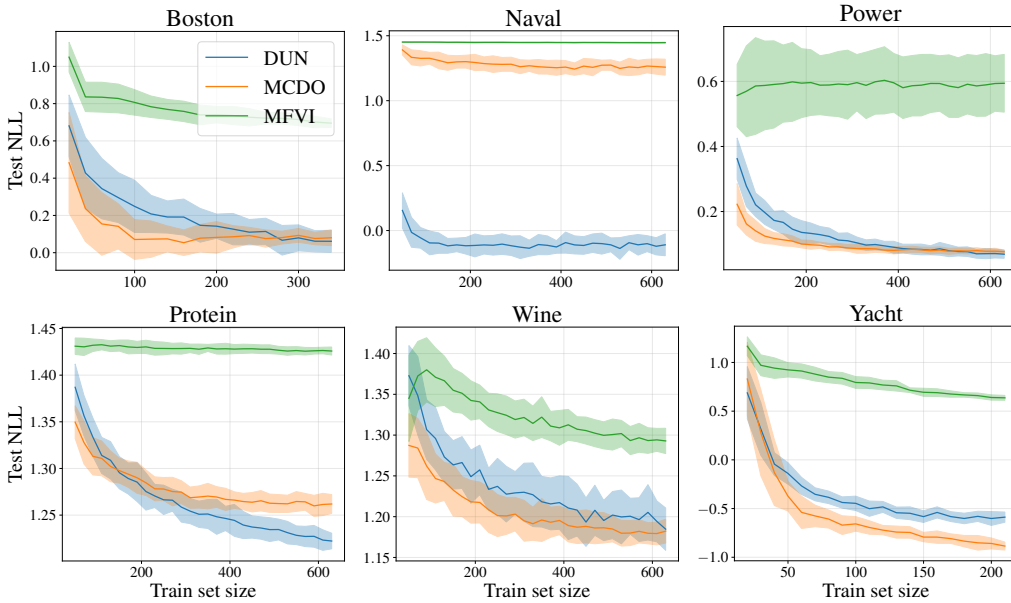


Figure D.2: Test NLL vs. number of training points evaluated on UCI datasets. **DUNs**, **MCDO** and **MFVI** are compared. Extension of the results in figure 3.

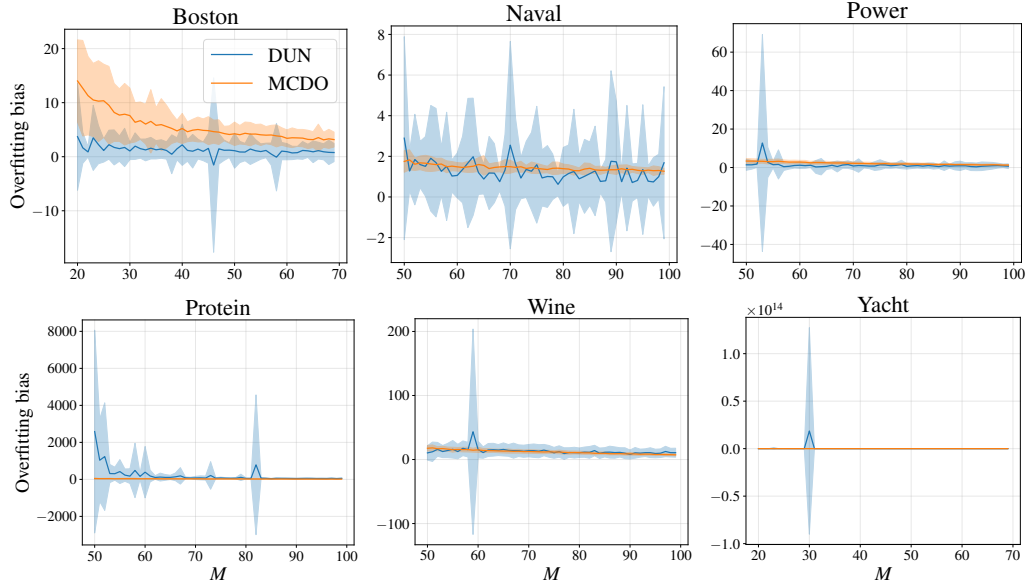


Figure D.3: Evolution of “overfitting bias”, as described by Farquhar et al. [2021]. See appendix B for more details. For **DUNs** and **MCDO** as the amount of acquired data M grows. DUNs are more robust to overfitting than MCDO. Extension of the results in figure 4.

D.2 Other datasets

The following two figures present results for MNIST (figure D.4) and the toy regression datasets (figure D.5). These results are consistent with the UCI regression results—DUNs either outperform the baseline methods (in MNIST, Foong et al. [2019], Matern and Wiggle), or perform comparably to MCDO (Simple 1d, Izmailov et al. [2020]).

D.2.1 MNIST results




Figure D.4: Test accuracy (left) and NLL (middle) vs. number of training points evaluated on MNIST. The performance of **DUNs** and **MCDO** is compared. Right: DUN posterior probabilities over depth, for the **smallest** and **largest** labelled datasets used in active learning.

D.2.2 Toy dataset results




Figure D.5: Test NLL vs. number of training points evaluated on toy datasets. The performance of **DUNs**, **MCDO** and **MFVI** is compared.

D.3 RMSE results

Figure D.6 shows the root mean squared error (RMSE) performance of DUNs, MCDO and MFVI on the UCI regression datasets. It is equivalent to figure 3 and figure D.2 but for RMSE rather than NLL. The RMSE results are consistent with the NLL results, with DUNs outperforming the baseline methods in five of the nine datasets, and performing comparably to MCDO in the remaining four datasets.




Figure D.6: Test RMSE vs. number of training points evaluated on UCI datasets. The performance of DUNs, MCDO and MFVI is compared.

D.4 Temperature of the proposal distribution

Figure D.7 shows the test NLL for DUNs using different temperatures T for the proposal distribution. The magnitude of T controls how deterministic the resulting sampling is—a larger T corresponds to more certainly selecting the point with the highest BALD score, while a smaller T is closer to uniform sampling. A temperature of $T = 10$ yields the best performance for most datasets.

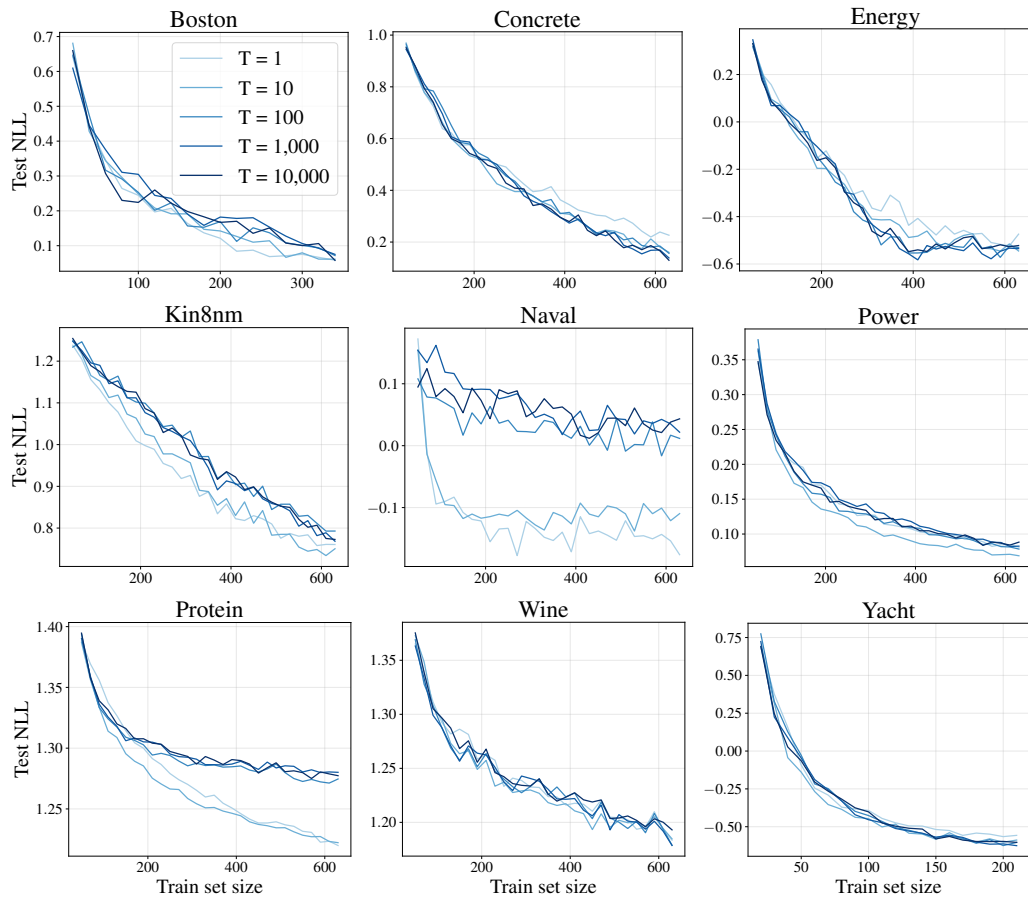


Figure D.7: Test NLL of DUNs using a stochastic relaxation of BALD. Different temperatures of the proposal distribution are compared. Means of 40 runs of the experiments are shown; standard deviations are not plotted for clarity.

D.5 Uniform and decaying depth priors

Figure D.8 compares DUNs' performance with a uniform prior over depth to that with exponentially decaying prior probabilities over depth. Much larger prior probabilities for earlier layers relative to later layers in the network should encourage a posterior that prefers shallower networks. This may, in turn, help to reduce overfitting in the early stages of active learning. In practice, no difference in performance is observed between the two prior distributions.

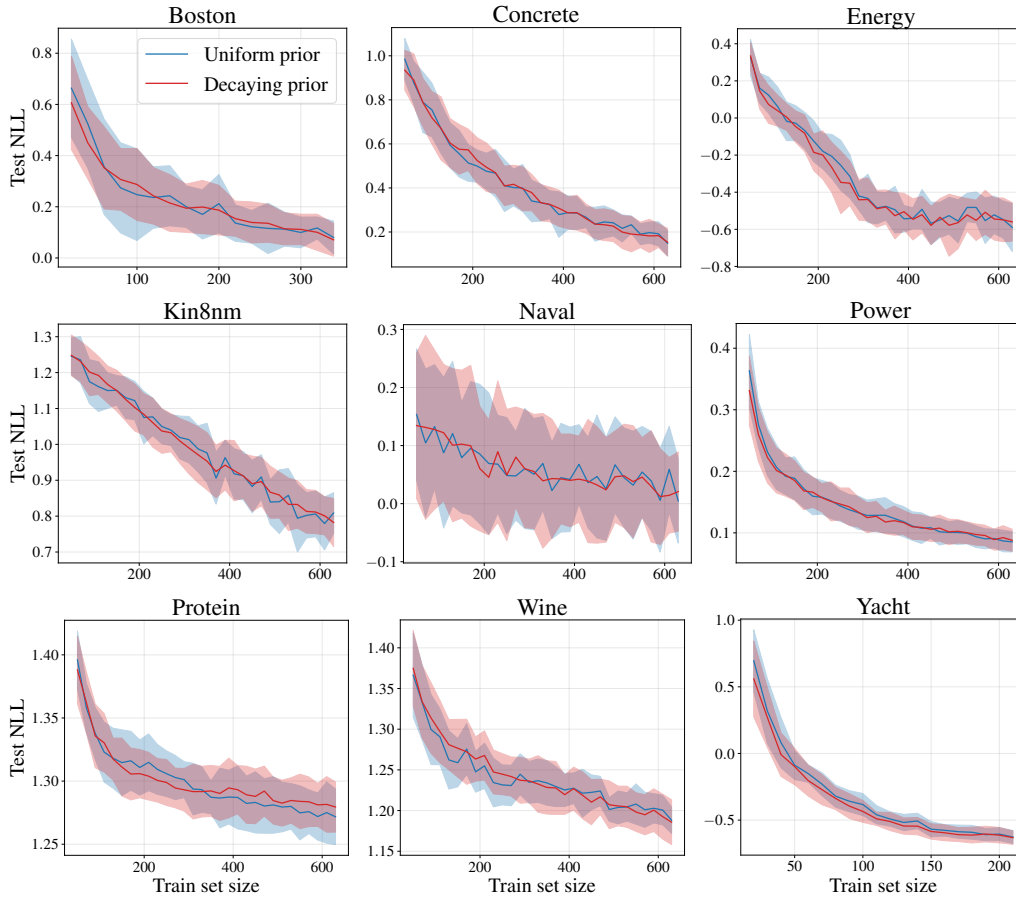



Figure D.8: Test NLL vs. number of training points for DUNs. A **uniform** depth prior is compared to a **decaying** prior.

D.6 Alternative depths for MCDO and MFVI

Figure D.9 compares the performance of DUNs to MCDO, MFVI and regular NNs trained with stochastic gradient descent (SGD), using different depth networks for the baseline methods. For each dataset and baseline method, a single hidden layer network and a network with depth equal to the highest posterior probability depth found by the DUN for that dataset are shown.



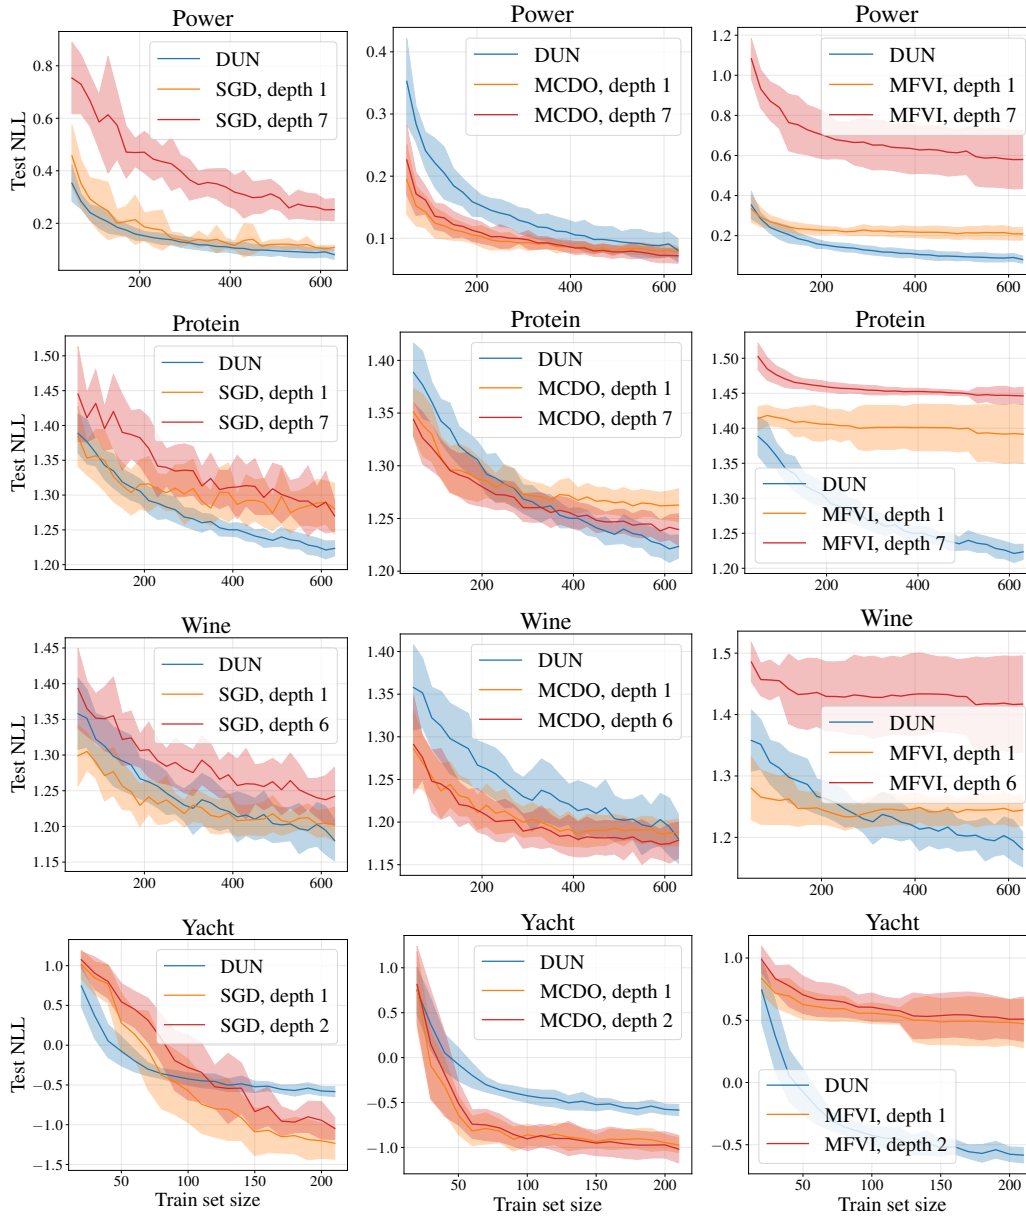


Figure D.9: NLL vs. number of training points evaluated on UCI datasets. The performance of SGD, MCDO and MFVI models with a single hidden layer and with the optimal number of layers found at the end of active learning by the DUN is compared to the performance of DUNs. A random acquisition strategy is used.