

---

# Addressing Bias in Active Learning with Depth Uncertainty Networks... or Not

---

Chelsea Murray James U. Allingham Javier Antorán José Miguel Hernández-Lobato  
Department of Engineering  
University of Cambridge  
{clm88, jua23, ja666, jmh233}@cam.ac.uk

## Abstract

Farquhar et al. [2021] show that correcting for active learning bias with underparameterised models leads to improved downstream performance. For overparameterised models such as NNs, however, correction leads either to decreased or unchanged performance. They suggest that this is due to an “overfitting bias” which offsets the active learning bias. We show that depth uncertainty networks operate in a low overfitting regime, much like underparameterised models. They should therefore see an increase in performance with bias correction. Surprisingly, they do not. We propose that this negative result, as well as the results Farquhar et al. [2021], can be explained via the lens of the bias-variance decomposition of generalisation error.

## 1 Introduction

Active learning improves data efficiency by identifying which examples maximise the expected gain in model performance [Cohn et al., 1995, Settles, 2010]. An acknowledged but relatively poorly understood problem with this approach, however, is that actively selecting informative training points introduces bias in inference, as the training data no longer follow the population distribution [MacKay, 1992]. Given that many statistical results rely on data points being identically and independently distributed (i.i.d.) samples from the population distribution, applying standard estimators to actively sampled datasets results in optimising for the wrong objective [Farquhar et al., 2021].

Recently, Farquhar et al. [2021] introduced an estimator,  $\tilde{R}_{\text{LURE}}$ , that corrects for the bias induced by active sampling by re-weighting the contribution of individual data points to the estimated population risk. Farquhar et al. [2021] find that, while underparameterised models such as linear regression benefit from the elimination of active learning bias, the performance of overparameterised models such as neural networks is either unaffected or negatively impacted by the unbiased risk estimator. This result is explained as active learning bias acting as a regulariser and helping to prevent flexible models from overfitting on small datasets.

Depth uncertainty networks (DUNs) are a class of Bayesian neural network (BNN) in which a prior is placed over the depth of the network, allowing us to favour simpler (i.e., shallower) functions. This enables DUNs to adapt their complexity to the varying size of the training dataset as additional labels are acquired during active learning. Because of this, we posit that DUNs will suffer less from overfitting in the early stages of training than other BNNs. This would imply that using the unbiased risk estimator of Farquhar et al. [2021] to eliminate active sampling bias should improve the performance of DUNs. Our main contributions are:

- Empirically validating the hypothesis that DUNs present less overfitting bias than other BNNs, specifically Monte Carlo dropout (MCDO), during active learning;

- Investigating the application of the  $\tilde{R}_{\text{LURE}}$  estimator to eliminate active learning bias from DUNs. Unexpectedly, despite DUNs having larger active learning bias than overfitting bias, we find that  $\tilde{R}_{\text{LURE}}$  does not result in improved downstream performance;
- Providing hypotheses for why we do not observe an impact on DUNs’ downstream performance when applying the unbiased risk estimator.

## 2 Background

### 2.1 Bias in active learning

We use notation similar to Farquhar et al. [2021]. Given a loss function  $\mathcal{L}(\mathbf{y}, f_{\theta}(\mathbf{x}))$ , we aim to find  $\theta$  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  $p_{\text{data}}$ . 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  $p_{\text{data}}$ . If these samples have been used to train  $\theta$ , 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:

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

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  $\tilde{R}_{\text{LURE}}$  (“levelled unbiased risk estimator”), which removes active learning bias:

$$\tilde{R}_{\text{LURE}} \equiv \frac{1}{M} \sum_{m=1}^M v_m \mathcal{L}_{i_m}; \quad v_m \equiv 1 + \frac{N - M}{N - m} \left( \frac{1}{(N - m + 1) \alpha(i_m; i_{1:m-1}, \mathcal{D}_{\text{pool}})} - 1 \right), \quad (2)$$

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  $\tilde{R}_{\text{LURE}}$  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.

### 2.2 Depth uncertainty networks for Active Learning

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  $\theta$  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 Bayesian active learning by disagreement (BALD) acquisition objective [Houlsby et al., 2011] are tractable and cheap to compute. See appendix A or [Antorán et al., 2020] for a detailed description.




Figure 1: DUN graphical model.

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 choose a prior that favours shallow models, limiting the flexibility of DUNs when few datapoints are observed. As the

actively acquired set increases, the model likelihood will start to dominate the prior, leading the posterior over depth to favour deeper models and thus more flexible functions. We show this in Figure C.6. We hypothesize that the capability to automatically adapt model flexibility as more data are observed will reduce overfitting in the small data regimes typical of active learning. Therefore DUNs, unlike standard BNNs, may in fact benefit from the use active learning bias corrective weights. It is worth noting that automatically adapting model complexity to the observed data is also a promised feature of traditional BNNs. However, unlike in DUNs where inferences are exact, weight space models require crude approximations which fall short in this regard [Ghosh et al., 2019].

### 3 Results

First, using the analysis of Farquhar et al. [2021], we show that the weighted estimator removes active learning bias for DUNs. We then show that DUNs overfit less than MCDO. Finally, we examine how training with  $\tilde{R}_{\text{LURE}}$  impacts downstream performance for DUNs. Experiments are performed on nine UCI regression datasets [Hernández-Lobato and Adams, 2015]; results for the Boston, Concrete and Energy datasets are presented in the following sections, with results for the rest provided in appendix C. The experimental setup is described in appendix B. All experiments are repeated 40 times, with the mean and standard deviations over the repetitions reported in our figures.

#### 3.1 Quantifying active learning bias

We estimate the extent of active learning bias when using  $\tilde{r}$  and verify the unbiasedness of  $\tilde{r}_{\text{LURE}}$ . A DUN is trained on 1,000 randomly selected data points from  $\mathcal{D}_{\text{pool}}$  using the standard negative log-likelihood (NLL) loss. We then estimate the model’s risk using both  $\tilde{r}$  and  $\tilde{r}_{\text{LURE}}$ , drawing  $M$  samples from an unobserved test set  $\mathcal{D}_{\text{test}}$ . The active learning bias is estimated as

$$B_{\text{ALB}}(\bullet) = r - E_{x_{1:M} \sim \alpha(x_{1:M}; \mathcal{D}_{\text{pool}})}[\bullet], \quad (3)$$

where  $\bullet$  is either  $\tilde{r}$  or  $\tilde{r}_{\text{LURE}}$  and  $r$  is computed on the whole  $\mathcal{D}_{\text{test}}$ . Figure 2 shows that  $\tilde{r}_{\text{LURE}}$  is unbiased, whereas the bias in  $\tilde{r}$  approaches zero only as  $M \rightarrow N$ .

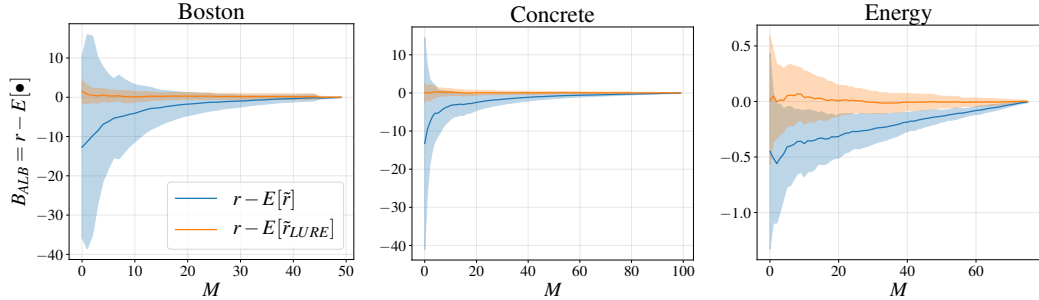


Figure 2: Active learning bias in risk estimation with DUNs, evaluated using  $\tilde{r}$  (blue) and  $\tilde{r}_{\text{LURE}}$  (orange). The unbiased risk estimator  $\tilde{r}_{\text{LURE}}$  eliminates active learning bias, while the standard estimator  $\tilde{r}$  tends to overestimate risk due to  $\alpha$  prioritising more difficult to fit points.

#### 3.2 Overfitting bias

To test the hypothesis that DUNs’ prior over depth reduces overfitting in low data regimes, we follow the procedure of Farquhar et al. [2021] to isolate overfitting bias. We train a model with  $M$  actively sampled points using  $\tilde{R}_{\text{LURE}}$  as the training objective. After acquiring each new point, we evaluate

$$B_{\text{OFB}} = r - \tilde{R}_{\text{LURE}}. \quad (4)$$

The difference between the risk evaluated on the test set ( $r$ ) and on the train set once the active learning bias is removed ( $\tilde{R}_{\text{LURE}}$ ) is taken as a measure of overfitting ( $B_{\text{OFB}}$ ).

Figure 3 shows that the magnitude of  $B_{\text{OFB}}$  for DUNs is small relative to BNNs trained with MCDO [Gal and Ghahramani, 2016], one of the methods used in Farquhar et al. [2021]. This suggests that DUNs are, indeed, robust to overfitting on small datasets. Overfitting bias in DUNs is also larger in magnitude than active learning bias for the datasets shown (comparing figure 3 to figure 2). This is not the case for MCDO (comparing figure 3 to figure C.2).

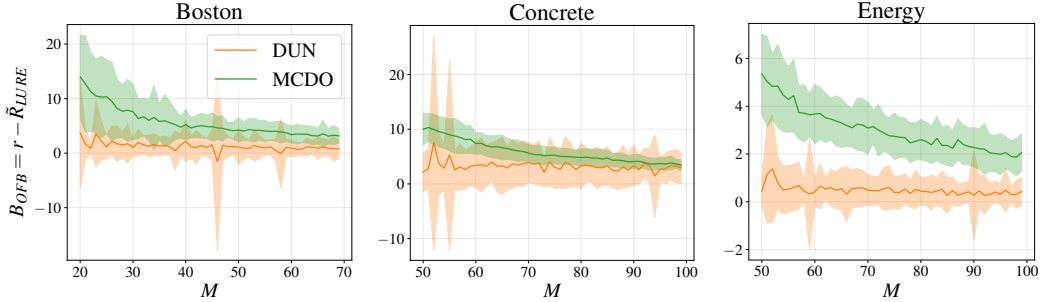


Figure 3: Evolution of overfitting bias for DUNs (orange) and MCDO (green) trained with  $\tilde{R}_{LURE}$  as the amount of acquired data  $M$  grows. Overfitting bias is smaller in DUNs than for MCDO.

### 3.3 Training with unbiased risk estimators

According to Farquhar et al. [2021], removing the active learning bias in cases where  $|B_{ALB}| \gg |B_{OFB}|$  should positively affect performance. Indeed, they observe this for linear models. Thus, having confirmed that the LURE weights eliminate active learning bias in DUNs, and that DUNs are in the  $|B_{ALB}| \gg |B_{OFB}|$  regime, at least for certain datasets, we evaluate whether using the unbiased estimator during training improves downstream performance. We compare the NLL obtained by evaluating our models on  $\mathcal{D}_{test}$  when using  $\tilde{R}$  and  $\tilde{R}_{LURE}$  as the training objectives. Surprisingly, figure 4 shows that using  $\tilde{R}_{LURE}$  does not affect NLL performance at all.




Figure 4: Test NLL as more training points are acquired for DUNs trained with  $\tilde{R}$  (blue) and  $\tilde{R}_{LURE}$  (orange). Active learning bias correction has minimal impact on downstream performance.

## 4 Discussion

In figure 4, we measure our models’ generalisation error, given by  $\text{bias}^2 + \text{variance}$  [Bishop, 2007]. Considering that neural networks are very flexible, we expect model bias to be negligible, and active learning bias to dominate. When we eliminate the latter with importance weights, as in equation (2), we pay the price of additional variance [Owen, 2013]. We therefore hypothesise that our results can be explained by gains due to bias reduction being lost due to increased variance. This increase in variance will depend on the choice of sampling distribution  $\alpha$ . “Overfitting bias”, as presented in equation (4) and [Farquhar et al., 2021] is computed using the train set and thus is not directly comparable with active learning bias in equation (3). Furthermore, the effect of “overfitting bias” is captured by model variance. The error induced by model variance is always additive with any statistical biases that may exist, precluding any cancellation effects. This hypothesis also explains why Farquhar et al. [2021] find  $\tilde{R}_{LURE}$  to be successful for linear models: since inflexible models are well-specified by small amounts of data, applying  $\tilde{R}_{LURE}$  barely increases variance. We also highlight that in the experimental design used by Farquhar et al. [2021] to obtain their linear model results, the temperature of the proposal distribution ( $T = 10,000$ ) is much larger than that used here ( $T = 10$ ). Higher temperatures approach a deterministic proposal, interfering with the assumption that  $\alpha(\cdot)$  has full support.

## References

- Javier Antorán, James Allingham, and José Miguel Hernández-Lobato. Depth uncertainty in neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.
- Christopher M. Bishop. *Pattern Recognition and Machine Learning, 5th Edition*. Information science and statistics. Springer, 2007.
- David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. Technical report, MASSACHUSETTS INST OF TECH CAMBRIDGE ARTIFICIAL INTELLIGENCE LAB, 1995.
- 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.
- 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>.
- 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.
- David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.
- Art B. Owen. *Monte Carlo theory, methods and examples*. 2013.
- 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.5. 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.5: 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.6 compares the graphical model representations of a BNN and a DUN. In BNNs, the weights  $\theta$  are treated as random variables drawn from a distribution  $p_\gamma(\theta)$  with hyperparameters  $\gamma$ . 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  $\beta$ :  $p_\beta(d) = \text{Cat}(d | \{\beta_i\}_{i=0}^D)$ . The model’s marginal log likelihood (MLL) is given by

$$\log p(\mathcal{D}_{\text{train}}; \theta) = \log \sum_{i=0}^D \left( p_\beta(d = i) \prod_{n=1}^N p(y_n | \mathbf{x}_n, d = i; \theta) \right), \quad (5)$$

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




Figure A.6: Left: graphical model of a BNN. Right: graphical model of a DUN.

The posterior,

$$p(d \mid \mathcal{D}_{\text{train}}; \theta) = \frac{p(\mathcal{D}_{\text{train}} \mid d; \theta)p_{\beta}(d)}{p(\mathcal{D}_{\text{train}}; \theta)}, \quad (6)$$

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 (5) with respect to  $\theta$  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  $\theta$  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:

$$\mathcal{L}(\phi, \theta) = \sum_{n=1}^N \mathbb{E}_{q_{\phi}(d)} [\log p(\mathbf{y}_n \mid \mathbf{x}_n, d; \theta)] - \text{KL}(q_{\phi}(d) \parallel p_{\beta}(d)), \quad (7)$$

allowing  $\theta$  and the variational parameters  $\phi$  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 (7) 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 (7) 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 Experimental setup

### B.1 Data

We implement batch-based active learning, with batches of 20 data points acquired in each query (with the exception of Yacht, which has a query size of 10 due to its smaller size). An initial training set representing a small proportion of the full data is selected uniformly at random in the first query. For all datasets, 80% of the data are used for training, 10% for validation and 10% for testing. Details about dataset sizes, input dimensionality, and active learning specifications for each dataset are provided in table B.1.

Table B.1: Summary of datasets and active learning specifications. 80% of the data are used for training, 10% for validation and 10% for testing.

NAME	SIZE	INPUT DIM.	INIT. TRAIN SIZE	NO. QUERIES	QUERY SIZE
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

### B.2 Model specifications

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. 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. For MCDO models a fixed dropout probability of 0.1 is used and prediction is based on 10 MC samples. DUNs use a uniform categorical prior over depth. The hyperparameter settings largely follow those used in Antorán et al. [2020] for their experiments with toy regression datasets.

All experiments are repeated 40 times with different weight initialisations and train-test splits (with the exception of the Protein dataset, for which experiments 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. Following Farquhar et al. [2021], the proposal distribution  $\alpha(i_m; i_{1:m-1}, \mathcal{D}_{\text{pool}})$  used is a stochastic relaxation of the Bayesian active learning by disagreement (BALD) acquisition function proposed by [Houlsby et al., 2011]. BALD scores are multiplied by a constant temperature  $T$  before being passed through the softmax function; candidate points are then sampled with probability proportional to the softmax probabilities. We use  $T = 10$  as this is found to yield the best performance for most of the datasets considered (results for experiments to determine the optimal value of  $T$  are given in appendix C.4).



## C Additional experimental results

### C.1 Quantifying active learning bias

Figure C.1 presents the results of the active learning bias experiments, as in figure 2, for the remaining datasets. Figure C.2 shows equivalent results for BNNs trained with MCDO.




Figure C.1: Bias introduced by actively sampling points with DUNs, evaluated using  $\tilde{R}$  (blue) and  $\tilde{R}_{LURE}$  (orange).

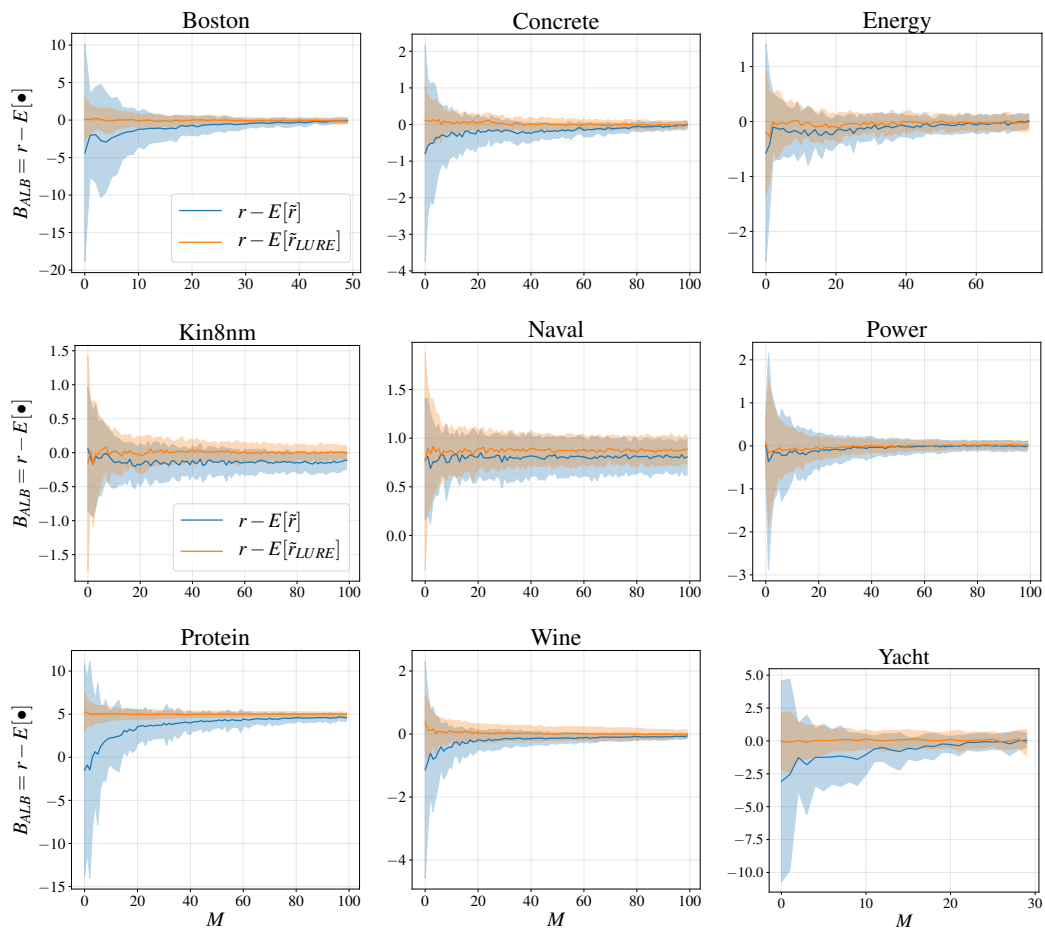


Figure C.2: Bias introduced by actively sampling points with MCDO, evaluated using  $\tilde{R}$  (blue) and  $\tilde{R}_{LURE}$  (orange).

## C.2 Overfitting bias

Figure C.3 compares the magnitude of overfitting bias for DUNs and BNNs trained with MCDO, as in figure 3. Figure C.3 shows additionally the overfitting bias when training with  $\tilde{R}$ .

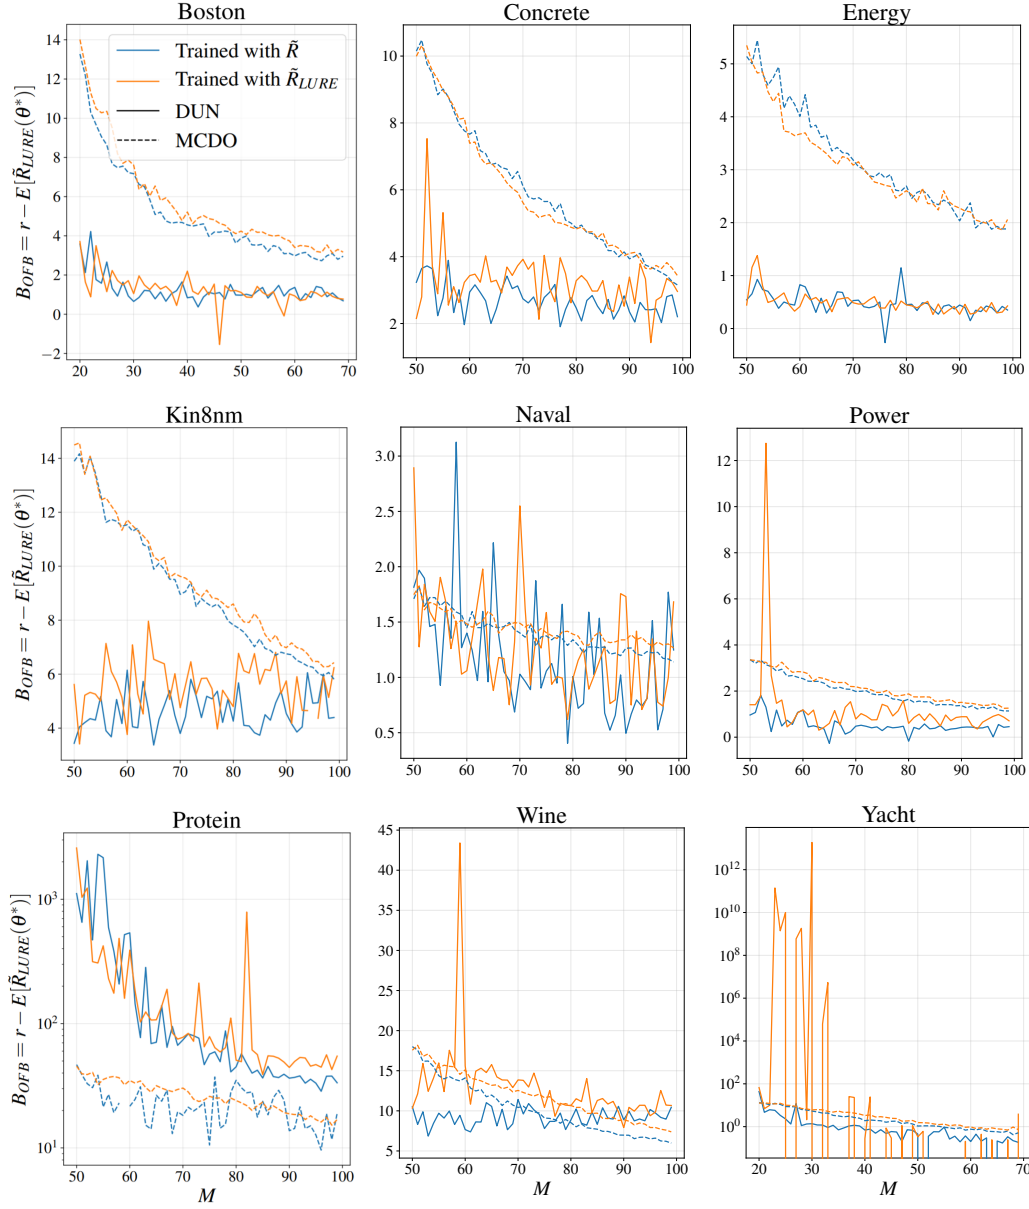


Figure C.3: Overfitting bias for DUNs (solid lines) and MCDO (dashed lines) trained with  $\tilde{R}$  (blue) and  $\tilde{R}_{LURE}$  (orange). NLL for Protein and Yacht is displayed in log scale.

### C.3 Training with unbiased risk estimators

Figure C.4 presents the results of training DUNs with  $\tilde{R}$  and  $\tilde{R}_{LURE}$ , as in figure 4, for the remaining datasets.

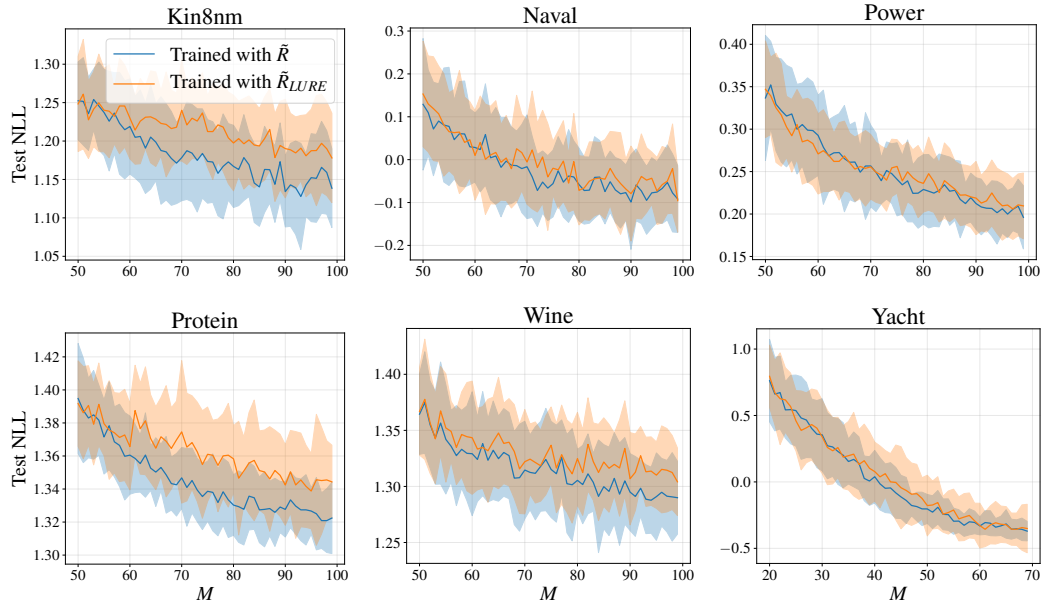


Figure C.4: Test NLL for DUNs trained with  $\tilde{R}$  (blue) and  $\tilde{R}_{LURE}$  (orange). A larger value of the orange line indicates that training with  $\tilde{R}_{LURE}$  harms performance, despite removing active learning bias.

#### C.4 Temperature of the proposal distribution

Figure C.5 shows the test NLL for DUNs using different temperatures  $T$  for the proposal distribution.

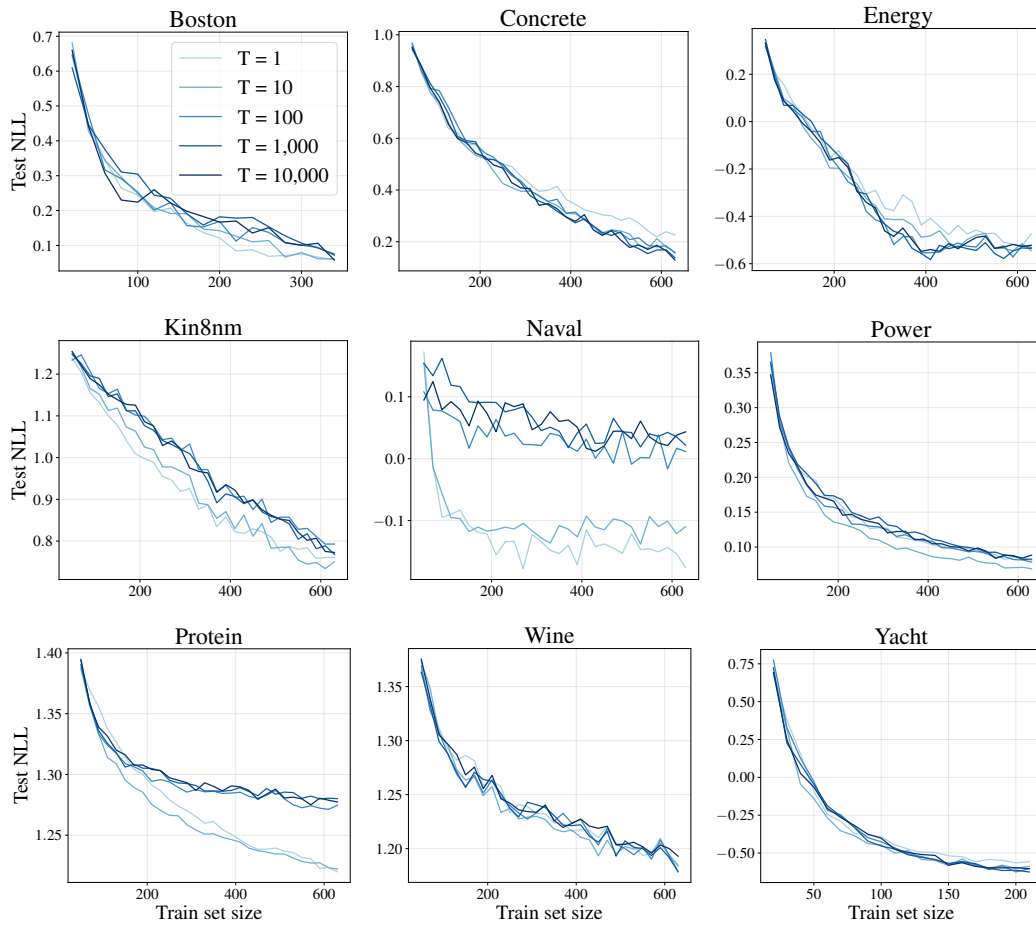


Figure C.5: Test NLL of DUNs using a stochastic relaxation of BALD. Different temperatures of the proposal distribution are compared. Means of 40 runs of the experiments are shown; standard deviations are not plotted for clarity.

### C.5 Depth posteriors for DUNs at different stages of active learning

Figure C.6 shows the posterior distributions over depth for DUNs for the smallest and largest dataset sizes used during active learning.

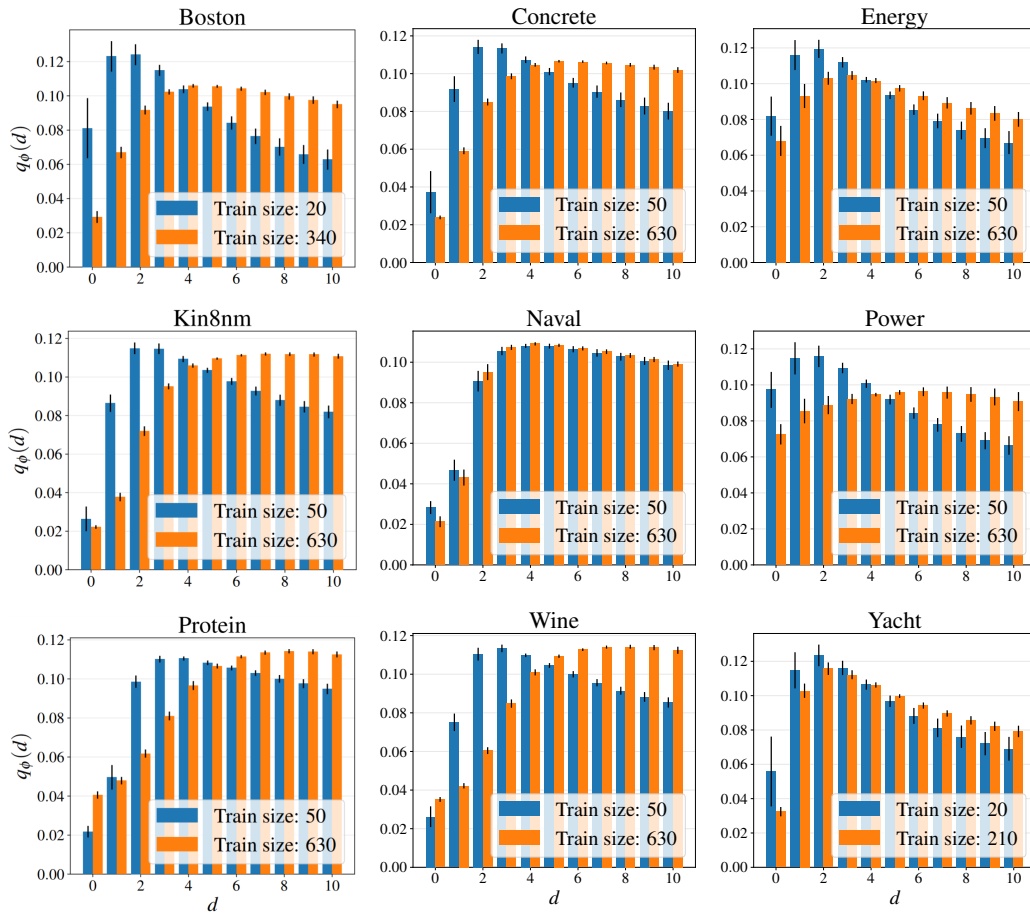


Figure C.6: Posterior probabilities over depth for DUNs trained on UCI regression datasets, for the smallest (blue bars) and largest (orange bars) labelled datasets used in active learning.