A   PROOFS FOR THE THEORETICAL RESULTS

We now provide the proofs for the results in Section 4.

A.1   PROOF OF PROPOSITION 1

Proof. Note that the posterior distributions $\widetilde{p}(\mathbf{w}|\boldsymbol{y}, \boldsymbol{\Phi})$ and $p_S(\mathbf{w}|\boldsymbol{y}, \boldsymbol{\Phi})$ are both Gaussian. We thus consider the squared 2-Wasserstein distance between two Gaussian distributions $\mathcal{N}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ and $\mathcal{N}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)$, which has the following closed-form expression (Givens et al., 1984)[4]:

$$W\left[\mathcal{N}\left(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1\right) \,||\, \mathcal{N}\left(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2\right)\right] = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + \mathrm{trace}\left(\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2 - 2\left(\boldsymbol{\Sigma}_1 \boldsymbol{\Sigma}_2\right)^{1/2}\right) . \quad (13)$$

Plugging in $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \mathbf{w}_{MAP}$, $\boldsymbol{\Sigma}_1 = H^{-1}$ and $\boldsymbol{\Sigma}_2 = \boldsymbol{M}_S \odot H^{-1}$, we obtain

$$\begin{aligned}
& W\left[\widetilde{p}(\mathbf{w}|\boldsymbol{y}, \boldsymbol{X}) \,||\, p_{\mathrm{S}}(\mathbf{w}|\boldsymbol{y}, \boldsymbol{X})\right] \\
&= W\left[\mathcal{N}(\mathbf{w}_{MAP}, H^{-1}) \,||\, \mathcal{N}(\mathbf{w}_{MAP}, \boldsymbol{M}_S \odot H^{-1})\right] \\
&= \|\mathbf{w}_{MAP} - \mathbf{w}_{MAP}\|_2^2 + \mathrm{trace}\left(H^{-1} + (\boldsymbol{M}_S \odot H^{-1}) - 2\left(H^{-1}(\boldsymbol{M}_S \odot H^{-1})\right)^{1/2}\right) \\
&= \mathrm{trace}\left((\mathbf{1} + \boldsymbol{M}_S) \odot H^{-1}\right) - \mathrm{trace}\left(2\left(H^{-1}\left(\boldsymbol{M}_S \odot H^{-1}\right)\right)^{1/2}\right) \\
&= \sum_{d=1}^{D} (1 + m_{dd})\,\sigma_d^2 - \mathrm{trace}\left(2\left(H^{-1}\left(\boldsymbol{M}_S \odot H^{-1}\right)\right)^{1/2}\right) \qquad \square
\end{aligned}$$

A.2   PROOF OF COROLLARY 1.1

Proof. For $H^{-1} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_D^2)$, the Wasserstein posterior gap in Eq. (12) simplifies to

$$W\left[\widetilde{p}(\mathbf{w}|\boldsymbol{y}, \boldsymbol{\Phi}) \,||\, p_{\mathrm{S}}(\mathbf{w}|\boldsymbol{y}, \boldsymbol{\Phi})\right] = \sum_{d=1}^{D} \left((1 + m_{dd})\,\sigma_d^2 - 2 m_{dd}\sigma_d^2\right) . \quad (14)$$

The optimal subnetwork selection strategy amounts to choosing the binary vector $\boldsymbol{m} = [m_{dd}]_{d=1}^{D}$ with $\sum_{d=1}^{D} m_d = S$ (i.e., we select S out of D parameters) s.t. the posterior gap in Eq. (14) is *minimized*. Observing that the contribution of the d-th parameter to the posterior gap is $(1+1)\sigma_d^2 - 1 \times 2\sigma_d^2 = 0$ if it is selected (i.e. if $m_{dd} = 1$), and $(1+0)\sigma_d^2 - 0 \times 2\sigma_d^2 = \sigma_d^2$ if it is *not* selected (i.e. if $m_{dd} = 0$), we see that the optimal subnetwork comprises of the S weights with the *largest* variances σ_d^2. $\square$

B   ADDITIONAL IMAGE CLASSIFICATION RESULTS

Table 1: AUC-ROC scores for out-of-distribution detection, using CIFAR10 vs SVHN and MNIST vs FashionMNIST as in- (source) and out-of-distribution (target) datasets, respectively (Nalisnick et al., 2019).

SOURCE	TARGET	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG
CIFAR10	SVHN	$0.85_{\pm 0.03}$	$0.86_{\pm 0.02}$	$0.85_{\pm 0.01}$	$0.86_{\pm 0.02}$	$0.91_{\pm 0.00}$	$0.86_{\pm 0.02}$	$0.83_{\pm 0.00}$
MNIST	Fashion	$0.92_{\pm 0.05}$	$0.75_{\pm 0.02}$	$0.82_{\pm 0.12}$	$0.75_{\pm 0.01}$	$0.90_{\pm 0.09}$	$0.72_{\pm 0.03}$	$0.97_{\pm 0.01}$

Table 2: MNIST – no rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.07_{\pm 0.01}$	$-0.01_{\pm 0.00}$	$-0.01_{\pm 0.00}$	$-0.04_{\pm 0.03}$	$-0.01_{\pm 0.00}$	$-0.01_{\pm 0.00}$	$-0.01_{\pm 0.00}$	$-0.14_{\pm nan}$
error	$0.01_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.01_{\pm 0.01}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.01_{\pm nan}$
ECE	$0.05_{\pm 0.01}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.10_{\pm nan}$
brier score	$0.02_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.02_{\pm 0.01}$	$0.01_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.04_{\pm nan}$

---

[4]This also holds for our case of a degenerate Gaussian with singular covariance matrix (Givens et al., 1984).
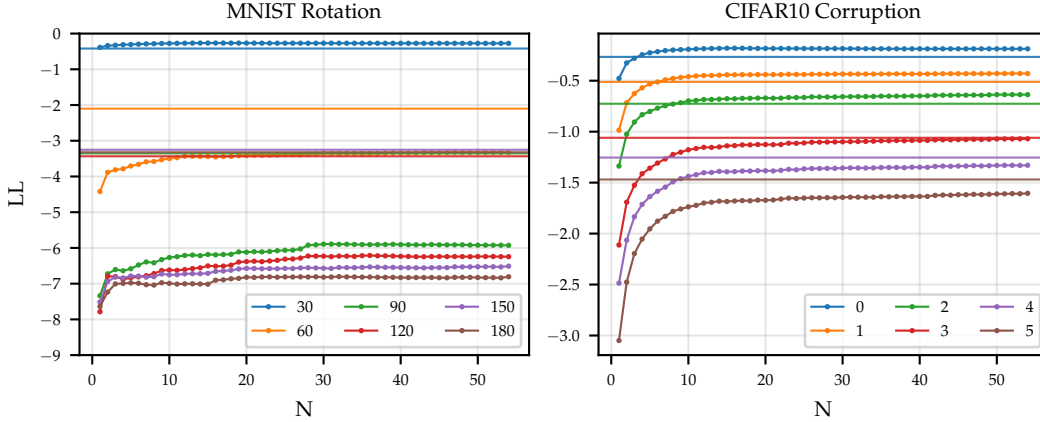
Figure 5: Rotated MNIST (left) and Corrupted CIFAR10 (right) results for deep ensembles (Lakshminarayanan et al., 2017) with large numbers of ensemble members (i.e. up to 55). Horizontal axis denotes number of ensemble members, and vertical axis denotes performance in terms of log-likelihood. Straight horizontal lines correspond to the performance of our method, as a reference. Colors denote different levels of rotation (left) and corruption (right). It can clearly be observed that the performance of deep ensembles saturates after around 15 ensemble members, meaning that adding more members yields strongly diminishing returns. This is in agreement with recent works (Antorán et al., 2020; Ashukha et al., 2020a; Lobacheva et al., 2020). Our method significantly outperforms even very large deep ensembles, especially for high degrees of rotation/corruption.


Figure 6: Rejection-classification plots. We simulate a realistic OOD rejection scenario (Filos et al., 2019b) by jointly evaluating our models on an in-distribution and an OOD test set. We allow our methods to reject increasing proportions of the data based on predictive entropy before classifying the rest. All predictions on OOD samples are treated as incorrect. Following (Nalisnick et al., 2019), we use CIFAR10 vs SVHN and MNIST vs FashionMNIST as in- and out-of-distribution datasets, respectively. Note that the SVHN test set is randomly sub-sampled down to a size of 10,000.
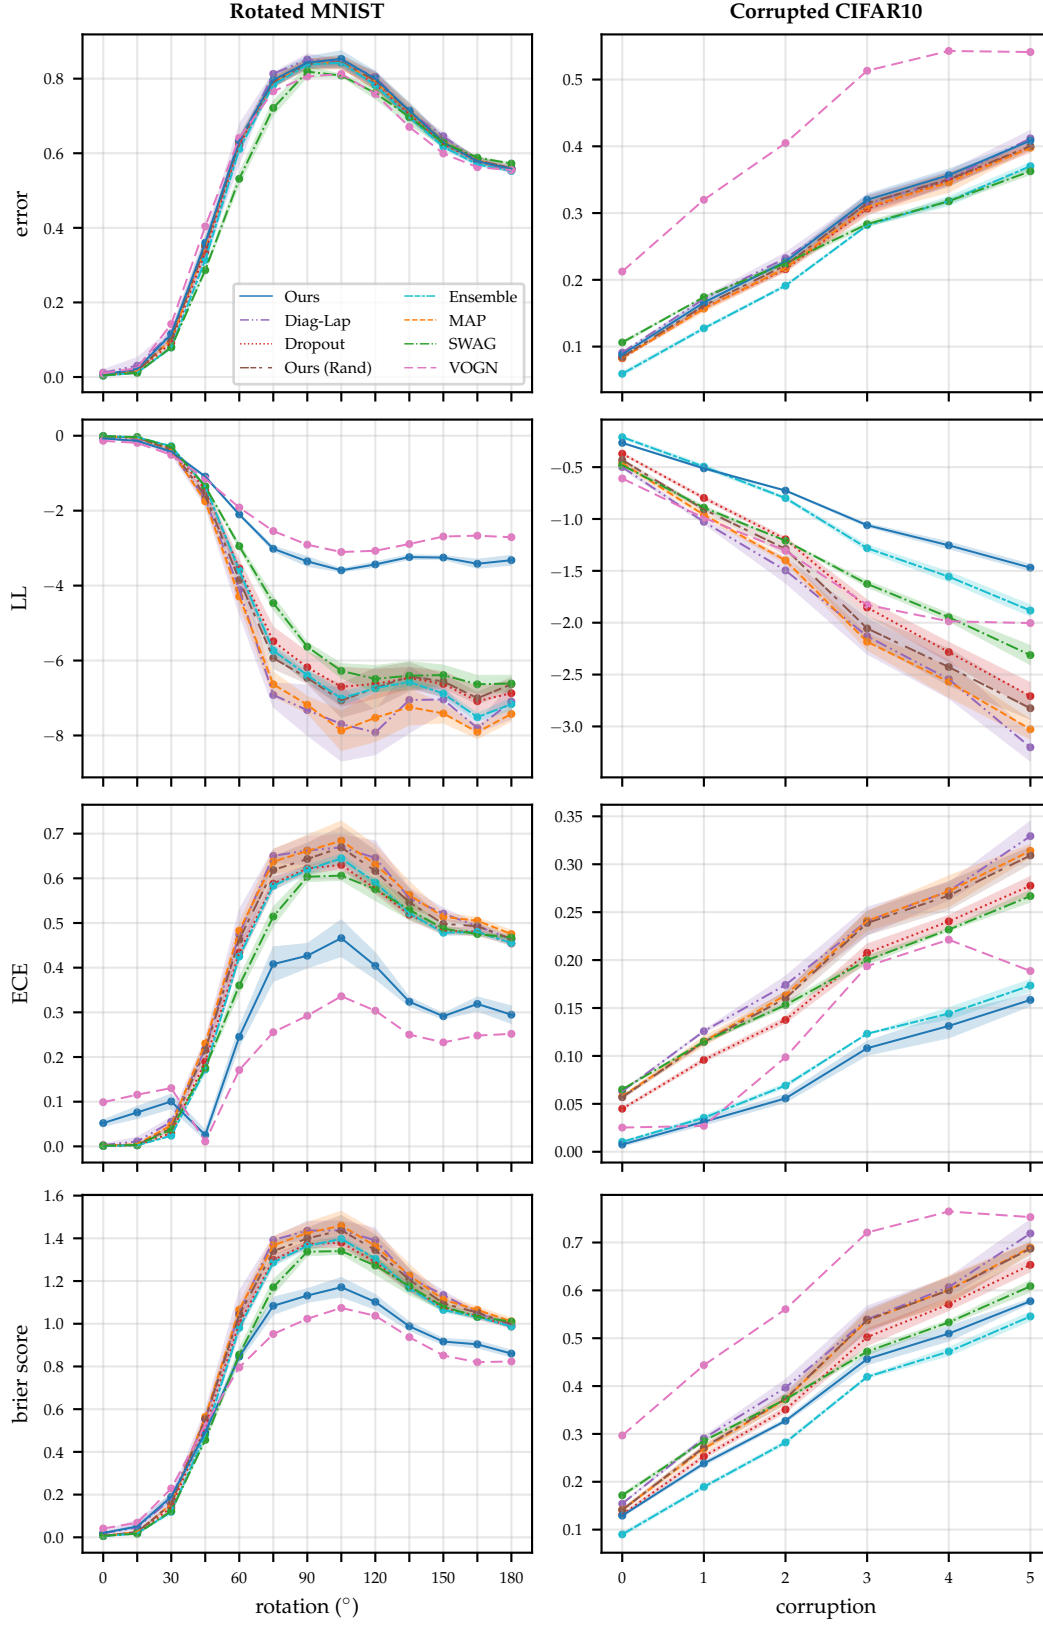
Figure 7: Full MNIST rotation and CIFAR10 corruption results, for ResNet-18, reporting predictive error, log-likelihood (LL), expected calibration error (ECE) and brier score, respectively (from top to bottom).

Figure 8: MNIST rotation results for ResNet-50, reporting predictive error, log-likelihood (LL), expected calibration error (ECE) and brier score. We choose a subnetwork containing only 0.167% (39,190 / 23,466,560) of the parameters of the full network. We see that subnetwork inference still results in an improvement in the calibration of predictive uncertainty. As expected, however, for ResNet-50 the improvement over MAP is smaller than for ResNet-18 where we were able to choose a subnetwork containing 0.38% of the parameters.

Table 3: MNIST – $15°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.14_{\pm 0.02}$	$-0.05_{\pm 0.00}$	$-0.05_{\pm 0.00}$	$-0.11_{\pm 0.08}$	$-0.04_{\pm 0.00}$	$-0.05_{\pm 0.00}$	$-0.04_{\pm 0.00}$	$-0.19_{\pm nan}$
error	$0.02_{\pm 0.00}$	$0.02_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.03_{\pm 0.02}$	$0.01_{\pm 0.00}$	$0.02_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.02_{\pm nan}$
ECE	$0.08_{\pm 0.01}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.01_{\pm 0.01}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.00_{\pm 0.00}$	$0.12_{\pm nan}$
brier score	$0.05_{\pm 0.01}$	$0.03_{\pm 0.00}$	$0.02_{\pm 0.00}$	$0.05_{\pm 0.03}$	$0.02_{\pm 0.00}$	$0.02_{\pm 0.00}$	$0.02_{\pm 0.00}$	$0.07_{\pm nan}$

Table 4: MNIST – $30°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.42_{\pm 0.04}$	$-0.36_{\pm 0.01}$	$-0.32_{\pm 0.02}$	$-0.44_{\pm 0.06}$	$-0.28_{\pm 0.02}$	$-0.39_{\pm 0.01}$	$-0.30_{\pm 0.00}$	$-0.51_{\pm nan}$
error	$0.11_{\pm 0.01}$	$0.10_{\pm 0.00}$	$0.09_{\pm 0.01}$	$0.12_{\pm 0.01}$	$0.08_{\pm 0.01}$	$0.10_{\pm 0.00}$	$0.08_{\pm 0.00}$	$0.14_{\pm nan}$
ECE	$0.10_{\pm 0.02}$	$0.04_{\pm 0.01}$	$0.03_{\pm 0.00}$	$0.06_{\pm 0.01}$	$0.02_{\pm 0.00}$	$0.05_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.13_{\pm nan}$
brier score	$0.19_{\pm 0.02}$	$0.16_{\pm 0.00}$	$0.14_{\pm 0.01}$	$0.18_{\pm 0.02}$	$0.12_{\pm 0.01}$	$0.16_{\pm 0.00}$	$0.12_{\pm 0.00}$	$0.23_{\pm nan}$

Table 5: MNIST – $45°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
midrule LL	$-1.09_{\pm 0.03}$	$-1.60_{\pm 0.05}$	$-1.44_{\pm 0.11}$	$-1.68_{\pm 0.20}$	$-1.36_{\pm 0.07}$	$-1.75_{\pm 0.06}$	$-1.35_{\pm 0.02}$	$-1.15_{\pm nan}$
error	$0.36_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.33_{\pm 0.01}$	$0.35_{\pm 0.03}$	$0.31_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.29_{\pm 0.00}$	$0.40_{\pm nan}$
ECE	$0.03_{\pm 0.01}$	$0.22_{\pm 0.01}$	$0.19_{\pm 0.02}$	$0.22_{\pm 0.02}$	$0.17_{\pm 0.01}$	$0.23_{\pm 0.01}$	$0.18_{\pm 0.00}$	$0.01_{\pm nan}$
brier score	$0.49_{\pm 0.02}$	$0.55_{\pm 0.02}$	$0.52_{\pm 0.02}$	$0.55_{\pm 0.04}$	$0.48_{\pm 0.02}$	$0.56_{\pm 0.02}$	$0.46_{\pm 0.01}$	$0.53_{\pm nan}$

Table 6: MNIST – $60°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-2.10_{\pm 0.03}$	$-3.85_{\pm 0.18}$	$-3.54_{\pm 0.23}$	$-4.11_{\pm 0.66}$	$-3.60_{\pm 0.10}$	$-4.29_{\pm 0.21}$	$-2.95_{\pm 0.08}$	$-1.92_{\pm nan}$
error	$0.63_{\pm 0.01}$	$0.63_{\pm 0.01}$	$0.62_{\pm 0.01}$	$0.62_{\pm 0.05}$	$0.61_{\pm 0.01}$	$0.63_{\pm 0.01}$	$0.53_{\pm 0.02}$	$0.64_{\pm nan}$
ECE	$0.25_{\pm 0.02}$	$0.46_{\pm 0.02}$	$0.43_{\pm 0.02}$	$0.47_{\pm 0.06}$	$0.42_{\pm 0.01}$	$0.48_{\pm 0.02}$	$0.36_{\pm 0.02}$	$0.17_{\pm nan}$
brier score	$0.85_{\pm 0.02}$	$1.04_{\pm 0.03}$	$1.00_{\pm 0.03}$	$1.05_{\pm 0.10}$	$0.98_{\pm 0.02}$	$1.07_{\pm 0.03}$	$0.86_{\pm 0.03}$	$0.80_{\pm nan}$

Table 7: MNIST – $75°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.02_{\pm 0.07}$	$-5.93_{\pm 0.28}$	$-5.49_{\pm 0.38}$	$-6.92_{\pm 0.32}$	$-5.74_{\pm 0.15}$	$-6.63_{\pm 0.33}$	$-4.46_{\pm 0.18}$	$-2.54_{\pm nan}$
error	$0.80_{\pm 0.02}$	$0.79_{\pm 0.01}$	$0.79_{\pm 0.01}$	$0.81_{\pm 0.00}$	$0.78_{\pm 0.01}$	$0.79_{\pm 0.01}$	$0.72_{\pm 0.02}$	$0.77_{\pm nan}$
ECE	$0.41_{\pm 0.04}$	$0.62_{\pm 0.03}$	$0.59_{\pm 0.01}$	$0.65_{\pm 0.01}$	$0.58_{\pm 0.01}$	$0.64_{\pm 0.03}$	$0.51_{\pm 0.02}$	$0.26_{\pm nan}$
brier score	$1.08_{\pm 0.04}$	$1.34_{\pm 0.04}$	$1.30_{\pm 0.02}$	$1.39_{\pm 0.01}$	$1.29_{\pm 0.02}$	$1.37_{\pm 0.04}$	$1.17_{\pm 0.04}$	$0.95_{\pm nan}$

Table 8: MNIST – $90°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.35_{\pm 0.13}$	$-6.46_{\pm 0.15}$	$-6.18_{\pm 0.41}$	$-7.32_{\pm 0.67}$	$-6.39_{\pm 0.17}$	$-7.18_{\pm 0.22}$	$-5.63_{\pm 0.12}$	$-2.91_{\pm nan}$
error	$0.84_{\pm 0.02}$	$0.84_{\pm 0.01}$	$0.84_{\pm 0.01}$	$0.85_{\pm 0.01}$	$0.84_{\pm 0.01}$	$0.84_{\pm 0.01}$	$0.82_{\pm 0.01}$	$0.81_{\pm nan}$
ECE	$0.43_{\pm 0.03}$	$0.64_{\pm 0.04}$	$0.62_{\pm 0.01}$	$0.66_{\pm 0.03}$	$0.62_{\pm 0.01}$	$0.66_{\pm 0.04}$	$0.60_{\pm 0.01}$	$0.29_{\pm nan}$
brier score	$1.13_{\pm 0.03}$	$1.40_{\pm 0.05}$	$1.37_{\pm 0.01}$	$1.44_{\pm 0.04}$	$1.36_{\pm 0.01}$	$1.43_{\pm 0.05}$	$1.34_{\pm 0.02}$	$1.02_{\pm nan}$

Table 9: MNIST – $105°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.59_{\pm 0.05}$	$-7.06_{\pm 0.45}$	$-6.70_{\pm 0.52}$	$-7.69_{\pm 0.99}$	$-7.01_{\pm 0.17}$	$-7.87_{\pm 0.53}$	$-6.28_{\pm 0.19}$	$-3.10_{\pm nan}$
error	$0.85_{\pm 0.02}$	$0.84_{\pm 0.02}$	$0.84_{\pm 0.01}$	$0.85_{\pm 0.01}$	$0.84_{\pm 0.01}$	$0.84_{\pm 0.02}$	$0.81_{\pm 0.00}$	$0.81_{\pm nan}$
ECE	$0.47_{\pm 0.04}$	$0.67_{\pm 0.05}$	$0.63_{\pm 0.01}$	$0.67_{\pm 0.03}$	$0.64_{\pm 0.01}$	$0.68_{\pm 0.04}$	$0.61_{\pm 0.01}$	$0.34_{\pm nan}$
brier score	$1.17_{\pm 0.05}$	$1.44_{\pm 0.07}$	$1.38_{\pm 0.02}$	$1.44_{\pm 0.04}$	$1.40_{\pm 0.01}$	$1.46_{\pm 0.07}$	$1.34_{\pm 0.02}$	$1.07_{\pm nan}$

Table 10: MNIST – $120°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.43_{\pm 0.07}$	$-6.73_{\pm 0.53}$	$-6.62_{\pm 0.39}$	$-7.92_{\pm 0.59}$	$-6.73_{\pm 0.11}$	$-7.53_{\pm 0.63}$	$-6.49_{\pm 0.36}$	$-3.07_{\pm nan}$
error	$0.80_{\pm 0.02}$	$0.79_{\pm 0.02}$	$0.78_{\pm 0.01}$	$0.81_{\pm 0.01}$	$0.78_{\pm 0.01}$	$0.79_{\pm 0.02}$	$0.76_{\pm 0.02}$	$0.76_{\pm nan}$
ECE	$0.40_{\pm 0.03}$	$0.62_{\pm 0.05}$	$0.58_{\pm 0.01}$	$0.65_{\pm 0.04}$	$0.59_{\pm 0.01}$	$0.63_{\pm 0.04}$	$0.58_{\pm 0.03}$	$0.30_{\pm nan}$
brier score	$1.10_{\pm 0.03}$	$1.35_{\pm 0.07}$	$1.29_{\pm 0.02}$	$1.39_{\pm 0.06}$	$1.30_{\pm 0.01}$	$1.36_{\pm 0.07}$	$1.27_{\pm 0.04}$	$1.04_{\pm nan}$

Table 11: MNIST – $135°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.24_{\pm 0.06}$	$-6.43_{\pm 0.38}$	$-6.46_{\pm 0.28}$	$-7.05_{\pm 0.88}$	$-6.57_{\pm 0.10}$	$-7.24_{\pm 0.48}$	$-6.40_{\pm 0.37}$	$-2.89_{\pm nan}$
error	$0.71_{\pm 0.02}$	$0.71_{\pm 0.02}$	$0.70_{\pm 0.01}$	$0.71_{\pm 0.01}$	$0.70_{\pm 0.01}$	$0.71_{\pm 0.02}$	$0.70_{\pm 0.02}$	$0.67_{\pm nan}$
ECE	$0.32_{\pm 0.01}$	$0.55_{\pm 0.03}$	$0.52_{\pm 0.01}$	$0.56_{\pm 0.02}$	$0.52_{\pm 0.01}$	$0.56_{\pm 0.03}$	$0.53_{\pm 0.02}$	$0.25_{\pm nan}$
brier score	$0.99_{\pm 0.02}$	$1.21_{\pm 0.05}$	$1.17_{\pm 0.02}$	$1.22_{\pm 0.04}$	$1.17_{\pm 0.01}$	$1.23_{\pm 0.05}$	$1.18_{\pm 0.04}$	$0.94_{\pm nan}$

Table 12: MNIST – $150°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.25_{\pm 0.05}$	$-6.56_{\pm 0.18}$	$-6.62_{\pm 0.33}$	$-7.04_{\pm 0.36}$	$-6.88_{\pm 0.11}$	$-7.41_{\pm 0.25}$	$-6.39_{\pm 0.27}$	$-2.69_{\pm nan}$
error	$0.63_{\pm 0.02}$	$0.63_{\pm 0.01}$	$0.63_{\pm 0.00}$	$0.65_{\pm 0.01}$	$0.62_{\pm 0.01}$	$0.63_{\pm 0.01}$	$0.63_{\pm 0.01}$	$0.60_{\pm nan}$
ECE	$0.29_{\pm 0.01}$	$0.50_{\pm 0.01}$	$0.48_{\pm 0.01}$	$0.52_{\pm 0.01}$	$0.48_{\pm 0.01}$	$0.51_{\pm 0.01}$	$0.49_{\pm 0.01}$	$0.23_{\pm nan}$
brier score	$0.92_{\pm 0.02}$	$1.10_{\pm 0.02}$	$1.07_{\pm 0.01}$	$1.13_{\pm 0.02}$	$1.06_{\pm 0.01}$	$1.11_{\pm 0.02}$	$1.08_{\pm 0.02}$	$0.85_{\pm nan}$

Table 13: MNIST – $165°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.42_{\pm 0.12}$	$-7.01_{\pm 0.15}$	$-7.08_{\pm 0.39}$	$-7.80_{\pm 0.12}$	$-7.51_{\pm 0.11}$	$-7.91_{\pm 0.18}$	$-6.63_{\pm 0.24}$	$-2.67_{\pm nan}$
error	$0.58_{\pm 0.01}$	$0.58_{\pm 0.01}$	$0.58_{\pm 0.01}$	$0.58_{\pm 0.00}$	$0.57_{\pm 0.01}$	$0.58_{\pm 0.01}$	$0.59_{\pm 0.00}$	$0.56_{\pm nan}$
ECE	$0.32_{\pm 0.02}$	$0.49_{\pm 0.01}$	$0.48_{\pm 0.01}$	$0.49_{\pm 0.01}$	$0.48_{\pm 0.00}$	$0.51_{\pm 0.01}$	$0.48_{\pm 0.00}$	$0.25_{\pm nan}$
brier score	$0.90_{\pm 0.02}$	$1.05_{\pm 0.01}$	$1.04_{\pm 0.01}$	$1.05_{\pm 0.01}$	$1.03_{\pm 0.01}$	$1.07_{\pm 0.02}$	$1.03_{\pm 0.01}$	$0.82_{\pm nan}$

Table 14: MNIST – $180°$ rotation.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-3.32_{\pm 0.13}$	$-6.63_{\pm 0.18}$	$-6.87_{\pm 0.32}$	$-7.10_{\pm 0.47}$	$-7.16_{\pm 0.16}$	$-7.43_{\pm 0.20}$	$-6.61_{\pm 0.22}$	$-2.71_{\pm nan}$
error	$0.56_{\pm 0.01}$	$0.56_{\pm 0.01}$	$0.56_{\pm 0.00}$	$0.55_{\pm 0.01}$	$0.55_{\pm 0.00}$	$0.56_{\pm 0.01}$	$0.57_{\pm 0.00}$	$0.55_{\pm nan}$
ECE	$0.29_{\pm 0.02}$	$0.46_{\pm 0.01}$	$0.45_{\pm 0.00}$	$0.46_{\pm 0.00}$	$0.46_{\pm 0.01}$	$0.48_{\pm 0.01}$	$0.47_{\pm 0.01}$	$0.25_{\pm nan}$
brier score	$0.86_{\pm 0.02}$	$1.00_{\pm 0.01}$	$0.99_{\pm 0.01}$	$0.99_{\pm 0.01}$	$0.99_{\pm 0.00}$	$1.01_{\pm 0.02}$	$1.01_{\pm 0.01}$	$0.82_{\pm nan}$

Table 15: CIFAR10 – no corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.27_{\pm 0.00}$	$-0.43_{\pm 0.01}$	$-0.37_{\pm 0.01}$	$-0.50_{\pm 0.02}$	$-0.21_{\pm 0.01}$	$-0.46_{\pm 0.02}$	$-0.48_{\pm 0.01}$	$-0.61_{\pm nan}$
error	$0.09_{\pm 0.00}$	$0.08_{\pm 0.00}$	$0.08_{\pm 0.00}$	$0.09_{\pm 0.00}$	$0.06_{\pm 0.00}$	$0.08_{\pm 0.00}$	$0.11_{\pm 0.00}$	$0.21_{\pm nan}$
ECE	$0.01_{\pm 0.00}$	$0.06_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.06_{\pm 0.00}$	$0.01_{\pm 0.00}$	$0.06_{\pm 0.00}$	$0.07_{\pm 0.00}$	$0.03_{\pm nan}$
brier score	$0.13_{\pm 0.00}$	$0.14_{\pm 0.00}$	$0.13_{\pm 0.00}$	$0.15_{\pm 0.00}$	$0.09_{\pm 0.00}$	$0.14_{\pm 0.00}$	$0.17_{\pm 0.00}$	$0.30_{\pm nan}$

Table 16: CIFAR10 – level 1 corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.51_{\pm 0.01}$	$-0.91_{\pm 0.01}$	$-0.80_{\pm 0.02}$	$-1.03_{\pm 0.02}$	$-0.50_{\pm 0.02}$	$-0.96_{\pm 0.02}$	$-0.89_{\pm 0.02}$	$-0.99_{\pm nan}$
error	$0.17_{\pm 0.01}$	$0.16_{\pm 0.00}$	$0.16_{\pm 0.00}$	$0.17_{\pm 0.00}$	$0.13_{\pm 0.00}$	$0.16_{\pm 0.00}$	$0.17_{\pm 0.00}$	$0.32_{\pm nan}$
ECE	$0.03_{\pm 0.00}$	$0.11_{\pm 0.00}$	$0.10_{\pm 0.00}$	$0.13_{\pm 0.00}$	$0.04_{\pm 0.00}$	$0.12_{\pm 0.01}$	$0.11_{\pm 0.00}$	$0.03_{\pm nan}$
brier score	$0.24_{\pm 0.00}$	$0.27_{\pm 0.00}$	$0.25_{\pm 0.00}$	$0.29_{\pm 0.00}$	$0.19_{\pm 0.00}$	$0.27_{\pm 0.01}$	$0.29_{\pm 0.00}$	$0.44_{\pm nan}$

Table 17: CIFAR10 – level 2 corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-0.73_{\pm 0.01}$	$-1.29_{\pm 0.06}$	$-1.20_{\pm 0.02}$	$-1.50_{\pm 0.12}$	$-0.80_{\pm 0.01}$	$-1.40_{\pm 0.03}$	$-1.21_{\pm 0.00}$	$-1.31_{\pm nan}$
error	$0.23_{\pm 0.00}$	$0.22_{\pm 0.01}$	$0.22_{\pm 0.00}$	$0.23_{\pm 0.01}$	$0.19_{\pm 0.00}$	$0.22_{\pm 0.00}$	$0.22_{\pm 0.00}$	$0.40_{\pm nan}$
ECE	$0.06_{\pm 0.00}$	$0.16_{\pm 0.01}$	$0.14_{\pm 0.00}$	$0.17_{\pm 0.01}$	$0.07_{\pm 0.00}$	$0.16_{\pm 0.00}$	$0.15_{\pm 0.00}$	$0.10_{\pm nan}$
brier score	$0.33_{\pm 0.00}$	$0.37_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.40_{\pm 0.02}$	$0.28_{\pm 0.00}$	$0.37_{\pm 0.01}$	$0.37_{\pm 0.00}$	$0.56_{\pm nan}$

Table 18: CIFAR10 – level 3 corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-1.06_{\pm 0.02}$	$-2.06_{\pm 0.12}$	$-1.85_{\pm 0.07}$	$-2.13_{\pm 0.17}$	$-1.28_{\pm 0.03}$	$-2.18_{\pm 0.08}$	$-1.63_{\pm 0.03}$	$-1.83_{\pm nan}$
error	$0.32_{\pm 0.01}$	$0.31_{\pm 0.01}$	$0.31_{\pm 0.01}$	$0.31_{\pm 0.01}$	$0.28_{\pm 0.00}$	$0.31_{\pm 0.01}$	$0.28_{\pm 0.00}$	$0.51_{\pm nan}$
ECE	$0.11_{\pm 0.01}$	$0.24_{\pm 0.01}$	$0.21_{\pm 0.01}$	$0.24_{\pm 0.01}$	$0.12_{\pm 0.00}$	$0.24_{\pm 0.01}$	$0.20_{\pm 0.00}$	$0.19_{\pm nan}$
brier score	$0.46_{\pm 0.01}$	$0.54_{\pm 0.02}$	$0.50_{\pm 0.02}$	$0.54_{\pm 0.03}$	$0.42_{\pm 0.00}$	$0.54_{\pm 0.02}$	$0.47_{\pm 0.01}$	$0.72_{\pm nan}$

Table 19: CIFAR10 – level 4 corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-1.25_{\pm 0.03}$	$-2.43_{\pm 0.18}$	$-2.28_{\pm 0.10}$	$-2.54_{\pm 0.18}$	$-1.56_{\pm 0.05}$	$-2.57_{\pm 0.15}$	$-1.95_{\pm 0.04}$	$-1.99_{\pm nan}$
error	$0.36_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.32_{\pm 0.01}$	$0.35_{\pm 0.01}$	$0.32_{\pm 0.00}$	$0.54_{\pm nan}$
ECE	$0.13_{\pm 0.01}$	$0.27_{\pm 0.01}$	$0.24_{\pm 0.01}$	$0.27_{\pm 0.01}$	$0.14_{\pm 0.01}$	$0.27_{\pm 0.02}$	$0.23_{\pm 0.00}$	$0.22_{\pm nan}$
brier score	$0.51_{\pm 0.02}$	$0.60_{\pm 0.03}$	$0.57_{\pm 0.01}$	$0.61_{\pm 0.02}$	$0.47_{\pm 0.01}$	$0.60_{\pm 0.03}$	$0.53_{\pm 0.00}$	$0.76_{\pm nan}$

Table 20: CIFAR10 – level 5 corruption.

	OURS	OURS (RAND)	DROPOUT	DIAG-LAP	ENSEMBLE	MAP	SWAG	VOGN
LL	$-1.47_{\pm 0.03}$	$-2.82_{\pm 0.11}$	$-2.71_{\pm 0.13}$	$-3.20_{\pm 0.13}$	$-1.88_{\pm 0.05}$	$-3.03_{\pm 0.10}$	$-2.31_{\pm 0.09}$	$-2.00_{\pm nan}$
error	$0.41_{\pm 0.00}$	$0.40_{\pm 0.01}$	$0.40_{\pm 0.01}$	$0.41_{\pm 0.01}$	$0.37_{\pm 0.01}$	$0.40_{\pm 0.00}$	$0.36_{\pm 0.01}$	$0.54_{\pm nan}$
ECE	$0.16_{\pm 0.01}$	$0.31_{\pm 0.01}$	$0.28_{\pm 0.01}$	$0.33_{\pm 0.02}$	$0.17_{\pm 0.01}$	$0.31_{\pm 0.01}$	$0.27_{\pm 0.01}$	$0.19_{\pm nan}$
brier score	$0.58_{\pm 0.00}$	$0.69_{\pm 0.01}$	$0.65_{\pm 0.01}$	$0.72_{\pm 0.03}$	$0.55_{\pm 0.01}$	$0.69_{\pm 0.01}$	$0.61_{\pm 0.01}$	$0.75_{\pm nan}$

C EXPERIMENTAL SETUP

C.1 TOY EXPERIMENTS

We train a single, 2 hidden layer network, with 50 hidden ReLU units per layer using MAP inference until convergence. Specifically, we use SGD with a learning rate of 1×10^{-3}, momentum of 0.9 and weight decay of 1×10^{-4}. We use a batch size of 512. The objective we optimise is the Gaussian log-likelihood of our data, where the mean is outputted by the network and the the variance is a hyperparameter learnt jointly with NN parameters by SGD. This variance parameters is shared among all datapoints. Once the network is trained, we perform post-hoc inference on it using different approaches. Since all of these involve the linearized approximation, the mean prediction is the same for all methods. Only their uncertainty estimates vary.

Note that while for this toy example, we could in principle use the full covariance matrix for the purpose of subnetwork selection, we still just use its diagonal (as described in Section 4) for consistency. We use GGN Laplace inference over network weights (not biases) in combination with the linearized predictive distribution in Eq. (8). Thus, all approaches considered share their predictive mean, allowing us to better compare their uncertainty estimates.

All approaches share a single prior precision of $\lambda = 3$. We chose to select the prior precision such that the full covariance approach (optimistic baseline) presents reasonable results. We use the same value for all other methods. We first tried a precision of 1 and found the full covariance approach to produce excessively large errorbars (covering the whole plot). A value of 3 produces more reasonable results.

Final layer inference is performed by computing the full Laplace covariance matrix and discarding all entries except those corresponding to the final layer of the NN. Results for random sub-network selection are obtained with a single sample from a scaled uniform distribution over weight choice.

C.2 UCI EXPERIMENTS

In this experiment, our fully connected NNs have numbers of hidden layers $h_d=\{1,2\}$ and hidden layer widths $w_d=\{50,100\}$. For a dataset with input dimension i_d, the number of weights is given by $D=(i_d+1)w_d+(h_d-1)w_d^2$. Our 2 hidden layer, 100 hidden unit models have a weight count of the order 10^4. The non-linearity used is ReLU.

We first obtain a MAP estimate of each model's weights. Specifically, we use SGD with a learning rate of 1×10^{-3}, momentum of 0.9 and weight decay of 1×10^{-4}. We use a batch size of 512. The objective we optimise is the Gaussian log-likelihood of our data, where the mean is outputted by the network and the the variance is a hyperparameter learnt jointly with NN parameters by SGD.

For each dataset split, we set aside 15% of the train data as a validation set. We use these for early stopping training. Training runs for a maximum of 2000 epochs but early stops with a patience of 500 if validation performance does not increase. For the larger Protein dataset, these values are 500 and 125. The weight settings which provide best validation performance are kept.

We then perform full network GGN Laplace inference for each model. We also use our proposed Wassertein rule together with the diagonal Hessian assumption to prune every network's weight variances such that the number of variances that remain matches the size of every smaller network under consideration. The prior precision used for these steps is chosen such that the resulting predictor's logliklihood performance on the validation set is maximised. Specifically, we employ a grid search over the values: $\lambda : [0.0001, 0.001, 0.1, 0.5, 1, 2, 5, 10, 100, 1000]$. In all cases, we employ the linearized predictive in Eq. (7). Consequently, networks with the same number of weights make the same mean predictions. Increasing the number of weight variances considered will thus only increase predictive uncertainty.

C.3 IMAGE EXPERIMENTS

The results shown in Section 5.3 and Appendix B are obtained by training ResNet-18 (and ResNet-50) models using SGD with momentum. For each experiment repetition, we train 7 different models: The first is for: 'MAP', 'Ours', 'Ours (Rand)', 'SWAG', 'Diag-Laplace' and as the first element of

'Ensemble'. We train 4 additional 'Ensemble' elements, 1 network with 'Dropout', and, finally 1 network for 'VOGN'. The methods 'Ours', 'Ours (Rand)', 'SWAG', and 'Diag-Laplace' are applied post training.

For all methods except 'VOGN' we use the following training procedure. The (initial) learning rate, momentum, and weight decay are 0.1, 0.9, and 1×10^{-4}, respectively. For 'MAP' we use 4 Nvidia P100 GPUs with a total batch size of 2048. For the calculation of the Jacobian in the subnetwork selection phase we use a single P100 GPU with a batch size of 4. For the calculation of the hessian we use a single P100 GPU with a batch size of 2. We train on 1 Nvidia P100 GPU with a batch size of 256 for all other methods. Each dataset is trained for a different number of epochs, shown in Table 21. We decay the learning rate by a factor of 10 at scheduled epochs, also shown in Table 21. Otherwise, all methods and datasets share hyperparameters. These hyperparameter settings are the defaults provided by `PyTorch` for training on ImageNet. We found them to perform well across the board. We report results obtained at the final training epoch. We do not use a separate validation set to determine the best epoch as we found ResNet-18 and ResNet-50 to not overfit with the chosen schedules.

Table 21: Per-dataset training configuration for image experiments.

DATASET	NO. EPOCHS	LR SCHEDULE
MNIST	90	40, 70
CIFAR10	300	150, 225

For 'Dropout', we add dropout to the standard ResNet-50 model (He et al., 2016) in between the 2 and 3 convolutions in the bottleneck blocks. This approach follows Zagoruyko & Komodakis (2016) and Ashukha et al. (2020b) who add dropout in-between the two convolutions of a WideResNet-50's basic block. Following Antorán et al. (2020), we choose a dropout probability of 0.1, as they found it to perform better than the value of 0.3 suggested by Ashukha et al. (2020b). We use 16 MC samples for predictions. 'Ensemble' uses 5 elements for prediction. Ensemble elements differ from each other in their initialisation, which is sampled from the He initialisation distribution (He et al., 2015). We do not use adversarial training as, inline with Ashukha et al. (2020b), we do not find it to improve results. For 'VOGN' we use the same procedure and hyper-parameters as used by Osawa et al. (2019) in their CIFAR10 experiments, with the exception that we use a learning rate of 1×10^{-3} as we we found a value of 1×10^{-4} not to result in convergence. We train on a single Nvidia P100 GPU with a batch size of 256. See the authors' GitHub for more details: `github.com/team-approx-bayes/dl-with-bayes/blob/master/distributed/classification/configs/cifar10/resnet18_vogn_bs256_8gpu.json`.

We modify the standard ResNet-50 and ResNet-18 architectures such that the first 7×7 convolution is replaced with a 3×3 convolution. Additionally, we remove the first max-pooling layer. Following Goyal et al. (2017), we zero-initialise the last batch normalisation layer in residual blocks so that they act as identity functions at the start of training.

At test time, we tune the prior precision used for 'Ours', 'Diag-Laplace' and 'SWAG' approximation on a validation set for each approach individually, as in Ritter et al. (2018); Kristiadi et al. (2020). We use a grid search from 1×10^{-4} to 1×10^{4} in logarithmic steps, and then a second, finer-grained grid search between the two best performing values (again with logarithmic steps).

C.4 DATASETS

The 1d toy dataset used in Section 5.1 was taken from Antorán et al. (2020). We obtained it from the authors' github repo: `https://github.com/cambridge-mlg/DUN`. Table 22 summarises the datasets used in Section 5.2. Wine and Protein are available from the UCI dataset repository Dua & Graff (2017). Kin8nm is available from `https://www.openml.org/d/189` Foong et al. (2019b). For the standard splits (Hernández-Lobato & Adams, 2015) 90% of the data is used for training and 10% for validation. For the gap splits (Foong et al., 2019b) a split is obtained per input dimension by ordering points by their values across that dimension and removing the middle 33% of

Table 22: Datasets from tabular regression used in Section 5.2

Dataset	N Train	N Val (15% train)	N Test	Splits	Output Dim	Output Type	Input Dim	Input Type
Wine	1223	216	160	20	1	Continous	11	Continous
Wine Gap	906	161	532	11	1	Continous	11	Continous
Kin8nm	6267	1106	819	20	1	Continous	8	Continous
Kin8nm Gap	4642	820	2730	8	1	Continous	8	Continous
Protein	34983	6174	4573	5	1	Continous	9	Continous
Protein Gap	25913	4573	15244	9	1	Continous	9	Continous

the points. These are used for validation. The datasets used for our image experiments are outlined in Table 23.

Table 23: Summary of image datasets. The test and train set sizes are shown in brackets, e.g. (test & train).

NAME	SIZE	INPUT DIM.	NO. CLASSES	NO. SPLITS
MNIST (LeCun et al., 1998)	70,000 (60,000 & 10,000)	784 (28 × 28)	10	2
Fashion-MNIST (Xiao et al., 2017)	70,000 (60,000 & 10,000)	784 (28 × 28)	10	2
CIFAR10 (Krizhevsky & Hinton, 2009)	60,000 (50,000 & 10,000)	3072 (32 × 32 × 3)	10	2
SVHN (Netzer et al., 2011)	99,289 (73,257 & 26,032)	3072 (32 × 32 × 3)	10	2