## Supplementary Material

Table of Contents
A Proofs \& Derivations ..... 15
A. 1 Function-Space Variational Objective ..... 15
A. 2 Distribution under Linearized Function Mapping ..... 16
B Further Empirical Results ..... 18
B. 1 Tabular Results for Diabetic Retinopathy Diagnosis Tasks ..... 18
B. 2 UCI Regression ..... 18
C Illustrative Examples ..... 19
C. 1 Two Moons Classification Task ..... 19
C. 2 Synthetic 1D Regression Datasets ..... 20
D Implementation, Training, and Evaluation Details ..... 21
D. 1 Hyperparameter Selection Protocol ..... 21
D. 2 FashionMNIST vs. MNIST/NotMNIST ..... 21
D. 3 CIFAR-10 vs. SVHN ..... 21
D. 4 Diabetic Retinopathy Diagnosis ..... 21
D. 5 Two Moons ..... 22
D. 6 1D Regression ..... 22
D. 7 Further Implementation Details ..... 22
D. 8 Selection of Context Distribution ..... 22
D. 9 Compute Resources ..... 22

## Appendix A Proofs \& Derivations

## A. 1 Function-Space Variational Objective

This proof follows steps from Matthews et al. [2016]. Consider measures $\hat{P}$ and $P$ both of which define distributions over some function $f$, indexed by an infinite index set $X$. Let $\mathcal{D}$ be a dataset and let $\mathbf{X}_{\mathcal{D}}$ denote a set of inputs and $\mathbf{y}_{\mathcal{D}}$ a set of targets. Consider the measure-theoretic version of Bayes' Theorem [Schervish, 1995]:

$$
\begin{equation*}
\frac{d \hat{P}}{d P}(f)=\frac{p_{X}(Y \mid f)}{p(Y)} \tag{A.1}
\end{equation*}
$$

where $p_{X}(Y \mid f)$ is the likelihood and $p(Y)=\int_{\mathbb{R}^{X}} p_{X}(Y \mid f) d P(f)$ is the marginal likelihood. We assume that the likelihood function is evaluated at a finite subset of the index set $X$. Denote by $\pi_{C}: \mathbb{R}^{X} \rightarrow \mathbb{R}^{C}$ a projection function that takes a function and returns the same function, evaluated at a finite set of points $C$, so we can write

$$
\begin{equation*}
\frac{d \hat{P}}{d P}(f)=\frac{d \hat{P}_{\mathbf{X}_{\mathcal{D}}}}{d P_{\mathbf{X}_{\mathcal{D}}}}\left(\pi_{\mathbf{X}_{\mathcal{D}}}(f)\right)=\frac{p\left(\mathbf{y}_{\mathcal{D}} \mid \pi_{\mathbf{X}_{\mathcal{D}}}(f)\right)}{p\left(\mathbf{y}_{\mathcal{D}}\right)} \tag{A.2}
\end{equation*}
$$

and similarly, the marginal likelihood becomes $p\left(\mathbf{y}_{\mathcal{D}}\right)=\int p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right) \mathrm{d} P_{\mathbf{X}_{\mathcal{D}}}\left(f_{\mathbf{X}_{\mathcal{D}}}\right)$. Now, considering the measure-theoretic version of the KL divergence between an approximating stochastic process $Q$ and a posterior stochastic process $\hat{P}$, we can write

$$
\begin{equation*}
\mathbb{D}_{\mathrm{KL}}(Q \| \hat{P})=\int \log \frac{d Q}{d P}(f) \mathrm{d} Q(f)-\int \log \frac{d \hat{P}}{d P}(f) \mathrm{d} Q(f) \tag{A.3}
\end{equation*}
$$

where $P$ is some prior stochastic process. Now, we can apply the measure-theoretic Bayes' Theorem to obtain

$$
\begin{align*}
\mathbb{D}_{\mathrm{KL}}(Q \| \hat{P}) & =\int \log \frac{d Q}{d P}(f) \mathrm{d} Q(f)-\int \log \frac{d \hat{P}}{d P}(f) \mathrm{d} Q(f)  \tag{A.4}\\
& =\int \log \frac{d Q^{\pi}}{d P^{\pi}}(f) \mathrm{d} Q^{\pi}(f)-\int \log \frac{d \hat{P}_{\mathbf{X}_{\mathcal{D}}}}{d P_{\mathbf{X}_{\mathcal{D}}}}\left(f_{\mathbf{X}_{\mathcal{D}}}\right) \mathrm{d} Q_{\mathbf{X}_{\mathcal{D}}}\left(f_{\mathbf{X}_{\mathcal{D}}}\right)  \tag{A.5}\\
& =\int \log \frac{d Q^{\pi}}{d P^{\pi}}(f) \mathrm{d} Q^{\pi}(f)-\mathbb{E}_{Q \mathbf{x}_{\mathcal{D}}}\left[\log p\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\log p\left(\mathbf{y}_{\mathcal{D}}\right), \tag{A.6}
\end{align*}
$$

where $\frac{d Q^{\pi}}{d P^{\pi}}(f)$ is marginally consistent given the projection $\pi$. Rearranging, we can get

$$
\begin{align*}
p\left(\mathbf{y}_{\mathcal{D}}\right) & =\mathbb{E}_{Q \mathbf{x}_{\mathcal{D}}}\left[\log p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\int \log \frac{d Q^{\pi}}{d P^{\pi}}(f) \mathrm{d} Q^{\pi}(f)+\mathbb{D}_{\mathrm{KL}}\left(Q^{\pi} \| \hat{P}\right)  \tag{A.7}\\
& \geq \mathbb{E}_{Q_{\mathbf{x}_{\mathcal{D}}}}\left[\log p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\int \log \frac{d Q^{\pi}}{d P^{\pi}}(f) \mathrm{d} Q^{\pi}(f)  \tag{A.8}\\
& =\mathbb{E}_{Q_{\mathbf{x}_{\mathcal{D}}}}\left[\log p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\mathbb{D}_{\mathrm{KL}}\left(Q^{\pi} \| P^{\pi}\right) . \tag{A.9}
\end{align*}
$$

Finally, this lower bound can equivalently be expressed as

$$
\begin{equation*}
p\left(\mathbf{y}_{\mathcal{D}}\right) \geq \mathbb{E}_{Q_{\mathbf{x}_{\mathcal{D}}}}\left[\log p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\mathbb{D}_{\mathrm{KL}}\left(Q_{\mathbf{x}_{\mathcal{D}}, \mathbf{x}_{\backslash \mathcal{D}}} \| P_{\mathbf{X}_{\mathcal{D}}, \mathbf{x}_{\backslash \mathcal{D}}}\right) \tag{A.10}
\end{equation*}
$$

where $\mathbf{X}_{\backslash \mathcal{D}}$ is an infinite index set excluding the finite index set $\mathbf{X}_{\mathcal{D}}$, that is, $\mathbf{X}_{\backslash \mathcal{D}} \cap \mathbf{X}_{\mathcal{D}}=\varnothing$, or by Theorem 1 in Sun et al. [2019], we can write

$$
\begin{equation*}
p\left(\mathbf{y}_{\mathcal{D}}\right) \geq \mathbb{E}_{Q_{\mathbf{x}_{\mathcal{D}}}}\left[\log p_{\mathbf{y} \mid f_{\mathbf{X}}}\left(\mathbf{y}_{\mathcal{D}} \mid f_{\mathbf{X}_{\mathcal{D}}}\right)\right]-\sup _{\mathbf{X} \in \mathcal{X}_{\mathbb{N}}} \mathbb{D}_{\mathrm{KL}}\left(Q_{\mathbf{X}} \| P_{\mathbf{X}}\right) \tag{A.11}
\end{equation*}
$$

where $\mathcal{X}_{\mathbb{N}} \doteq \bigcup_{n \in \mathbb{N}}\left\{\mathbf{X} \in \mathcal{X}_{n} \mid \mathcal{X}_{n} \subseteq \mathbb{R}^{n \times D}\right\}$ is the collection of all finite sets of evaluation points.

## A. 2 Distribution under Linearized Function Mapping

Proposition 1 (Distribution under Linearized Mapping). For a stochastic function $f(\cdot ; \boldsymbol{\Theta})$ defined in terms of stochastic parameters $\boldsymbol{\Theta}$ distributed according to distribution $g_{\boldsymbol{\Theta}}$ with $\mathbf{m} \doteq \mathbb{E}_{g_{\Theta}}[\boldsymbol{\Theta}]$ and $\mathbf{S} \doteq \operatorname{Cov}_{g_{\boldsymbol{\Theta}}}[\boldsymbol{\Theta}]$, denote the linearization of the stochastic function $f(\cdot ; \boldsymbol{\Theta})$ about $\mathbf{m}$ by

$$
f(\cdot ; \boldsymbol{\Theta}) \approx \tilde{f}(\cdot ; \mathbf{m}, \boldsymbol{\Theta}) \doteq f(\cdot ; \mathbf{m})+\mathcal{J}(\cdot ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m})
$$

where $\left.\mathcal{J}(\cdot ; \mathbf{m}) \doteq(\partial f(\cdot ; \boldsymbol{\Theta}) / \partial \boldsymbol{\Theta})\right|_{\boldsymbol{\Theta}=\mathbf{m}}$ is the Jacobian of $f(\cdot ; \boldsymbol{\Theta})$ evaluated at $\boldsymbol{\Theta}=\mathbf{m}$. Then the mean and co-variance of the distribution over the linearized mapping $f$ at $\mathbf{X}, \mathbf{X}^{\prime} \in \mathcal{X}$ are given by $\mathbb{E}[\tilde{f}(\mathbf{X} ; \boldsymbol{\Theta})]=f(\mathbf{X} ; \mathbf{m})$

$$
\operatorname{Cov}\left[\tilde{f}(\mathbf{X} ; \boldsymbol{\Theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \boldsymbol{\Theta}\right)\right]=\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbf{S} \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}
$$

Proof. We wish to find $\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})]$ and

$$
\begin{align*}
& \operatorname{Cov}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\theta}\right)\right) \\
& =\mathbb{E}\left[(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})-\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})])\left(\tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\theta}\right)-\mathbb{E}\left[\tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\theta}\right)\right]\right)^{\top}\right] . \tag{A.12}
\end{align*}
$$

To see that $\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})]=f(\mathbf{X} ; \mathbf{m})$, note that, by linearity of expectation, we have

$$
\begin{align*}
\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})] & =\mathbb{E}[f(\mathbf{X} ; \mathbf{m})+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\mathbf{\Theta}-\mathbf{m})] \\
& =f(\mathbf{X} ; \mathbf{m})+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\mathbb{E}[\mathbf{\Theta}]-\mathbf{m})=f(\mathbf{X} ; \mathbf{m}) \tag{A.13}
\end{align*}
$$

To see that $\operatorname{Cov}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\theta}\right)\right)=\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbf{S} \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}$, note that in general, for a multivariate random variable $\mathbf{Z}, \operatorname{Cov}(\mathbf{Z}, \mathbf{Z})=\mathbb{E}\left[\mathbf{Z} \mathbf{Z}^{\top}\right]+\mathbb{E}[\mathbf{Z}] \mathbb{E}[\mathbf{Z}]^{\top}$, and hence,

$$
\begin{align*}
& \operatorname{Cov}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)\right) \\
& =\mathbb{E}\left[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)^{\top}\right]-\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})] \mathbb{E}\left[\tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)\right]^{\top} \tag{A.14}
\end{align*}
$$

We already know that $\mathbb{E}[\tilde{f}(\mathbf{X} ; \boldsymbol{\Theta})]=f(\mathbf{X} ; \mathbf{m})$, so we only need to find $\mathbb{E}\left[\tilde{f}(\mathbf{X} ; \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{\Theta}\right)^{\top}\right]$ :

$$
\begin{align*}
& \mathbb{E}_{g_{\boldsymbol{\Theta}}}\left[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)^{\top}\right] \\
& =\mathbb{E}_{g_{\Theta}}\left[(f(\mathbf{X} ; \mathbf{m})+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m}))\left(f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)+\mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)(\boldsymbol{\Theta}-\mathbf{m})\right)^{\top}\right]  \tag{A.15}\\
& =\mathbb{E}_{g_{\boldsymbol{\Theta}}}\left[f(\mathbf{X} ; \mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}+(\mathcal{J}(\mathbf{X} ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m}))\left(\mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)(\boldsymbol{\Theta}-\mathbf{m})\right)^{\top}\right.  \tag{A.16}\\
& \left.\quad+f(\mathbf{X} ; \mathbf{m})\left(\mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)(\boldsymbol{\Theta}-\mathbf{m})\right)^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}\right] \\
& =\mathbb{E}_{g_{\Theta}}\left[f(\mathbf{X} ; \mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m})(\mathbf{\Theta}-\mathbf{m})^{\top} \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}\right. \\
& \left.\quad+f(\mathbf{X} ; \mathbf{m})\left(\mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)(\boldsymbol{\Theta}-\mathbf{m})\right)^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}\right]  \tag{A.17}\\
& =f(\mathbf{X} ; \mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbb{E}_{g_{\Theta}}\left[(\boldsymbol{\Theta}-\mathbf{m})(\mathbf{\Theta}-\mathbf{m})^{\top}\right] \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top} \\
& \quad+f(\mathbf{X} ; \mathbf{m})(\mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)(\underbrace{\left.\mathbb{E}_{g_{\Theta}}[\boldsymbol{\Theta}]-\mathbf{m}\right)}_{=0})^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m})(\underbrace{\left.\mathbb{E}_{g_{\Theta}}[\boldsymbol{\Theta}]-\mathbf{m}\right)}_{=0}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top} \tag{A.18}
\end{align*}
$$

where the last line follows from the definition of $g_{\Theta}$. By definition of the covariance, we then obtain

$$
\begin{align*}
& \mathbb{E}_{g_{\boldsymbol{\Theta}}}\left[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)^{\top}\right] \\
& =f(\mathbf{X} ; \mathbf{m}) f\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbb{E}_{g_{\boldsymbol{\Theta}}}\left[(\boldsymbol{\Theta}-\mathbf{m})(\boldsymbol{\Theta}-\mathbf{m})^{\top}\right] \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}  \tag{A.19}\\
& =f(\mathbf{X} ; \mathbf{m}) f(\mathbf{X} ; \mathbf{m})^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m}) \operatorname{Cov}(\boldsymbol{\Theta}) \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top} \tag{A.20}
\end{align*}
$$

With this result, we obtain the covariance function

$$
\begin{align*}
& \operatorname{Cov}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)\right) \\
& =\mathbb{E}\left[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)^{\top}\right]-\mathbb{E}[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})] \mathbb{E}\left[\tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)\right]^{\top}  \tag{A.21}\\
& =\mathbb{E}\left[\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}) \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)^{\top}\right]-f(\mathbf{X} ; \mathbf{m}) f(\mathbf{X} ; \mathbf{m})^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m}) \operatorname{Cov}(\mathbf{\Theta}) \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}  \tag{A.22}\\
& \left.=f(\mathbf{X} ; \boldsymbol{\Theta}) f\left(\mathbf{X}^{\prime} ; \boldsymbol{\Theta}\right)^{\top}-f(\mathbf{X} ; \mathbf{m}) f(\mathbf{X} ; \mathbf{m})^{\top}+\mathcal{J}(\mathbf{X} ; \mathbf{m}) \operatorname{Cov} \boldsymbol{\Theta}\right) \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}  \tag{A.23}\\
& =\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbb{V}[\boldsymbol{\Theta}] \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top} \tag{A.24}
\end{align*}
$$

Finally, $\operatorname{Cov}(\boldsymbol{\Theta})=\mathbf{S}$ yields $\operatorname{Cov}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}), \tilde{f}\left(\mathbf{X}^{\prime} ; \mathbf{m}, \boldsymbol{\Theta}\right)\right)=\mathcal{J}(\mathbf{X} ; \mathbf{m}) \mathbf{S} \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)^{\top}$. This concludes the proof.

Proposition 2 (Approximate Distribution under Linearized Mapping). For a stochastic function $f(\cdot ; \boldsymbol{\Theta})$ defined in terms of stochastic parameters $\boldsymbol{\Theta}$ distributed according to distribution $g_{\boldsymbol{\Theta}}=\mathcal{N}(\mathbf{m}, \mathbf{S})$, denote the linearization of the stochastic function $f(\cdot ; \boldsymbol{\Theta})$ about $\mathbf{m}$ by

$$
f(\cdot ; \boldsymbol{\Theta}) \approx \tilde{f}(\cdot ; \mathbf{m}, \boldsymbol{\Theta}) \doteq f(\cdot ; \mathbf{m})+\mathcal{J}(\cdot ; \mathbf{m})(\boldsymbol{\Theta}-\mathbf{m})
$$

where $\left.\mathcal{J}(\cdot ; \mathbf{m}) \doteq(\partial f(\cdot ; \boldsymbol{\Theta}) / \partial \boldsymbol{\Theta})\right|_{\boldsymbol{\Theta}=\mathbf{m}}$ is the Jacobian of $f(\cdot ; \boldsymbol{\Theta})$ evaluated at $\boldsymbol{\Theta}=\mathbf{m}$. Then, for a partition of the set of parameters into sets $\alpha$ and $\beta$, a distribution $g_{\Theta}=\mathcal{N}(\mathbf{m}, \mathbf{S})$ with $\boldsymbol{\Theta}_{\alpha} \perp \boldsymbol{\Theta}_{\beta}$, the distribution $\tilde{g}_{\tilde{f}(\mathbf{X} ; \Theta)}$ can be approximated via the Monte Carlo estimator

$$
\begin{equation*}
\hat{\tilde{g}}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}=\frac{1}{R} \sum_{j=1}^{R} \mathcal{N}\left(f(\mathbf{X} ; \mathbf{m})+\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)^{(j)}, \mathcal{J}_{\beta}(\mathbf{X} ; \mathbf{m}) \mathbf{S}_{\beta} \mathcal{J}\left(\mathbf{X}^{\prime} ; \mathbf{m}\right)_{\beta}^{\top}\right) \tag{A.25}
\end{equation*}
$$

where $g_{\boldsymbol{\Theta}_{\alpha}}=\mathcal{N}\left(\mathbf{m}_{\alpha}, \mathbf{S}_{\alpha}\right), g_{\boldsymbol{\Theta}_{\beta}}=\mathcal{N}\left(\mathbf{m}_{\beta}, \mathbf{S}_{\beta}\right)$, and

$$
\begin{equation*}
\tilde{f}_{\alpha}\left(\cdot ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right) \doteq \mathcal{J}_{\alpha}(\cdot ; \mathbf{m})\left(\boldsymbol{\Theta}_{\alpha}-\mathbf{m}_{\alpha}\right) \tag{A.26}
\end{equation*}
$$

with $\mathcal{J}_{\alpha}(\cdot ; \mathbf{m})$ denoting the columns of the Jacobian matrix corresponding to the sets of parameters $\alpha$ and $\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \mathbf{\Theta}_{\alpha}\right)^{(j)}$ for $j=1, \ldots, R$ obtained by sampling parameters from the distribution $g_{\boldsymbol{\Theta}_{\alpha}}=\mathcal{N}\left(\mathbf{m}_{\alpha}, \mathbf{S}_{\alpha}\right)$.

Proof. Consider a partition of the set of parameters into sets $\alpha$ and $\beta$ and express the linearized mapping as

$$
\begin{equation*}
\tilde{f}(\cdot ; \mathbf{m}, \boldsymbol{\Theta})=\tilde{f}_{\alpha}\left(\cdot ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)+\tilde{f}_{\beta}\left(\cdot ; \mathbf{m}, \boldsymbol{\Theta}_{\beta}\right) \tag{A.27}
\end{equation*}
$$

with

$$
\begin{equation*}
\tilde{f}_{\alpha}\left(\cdot ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right) \doteq \mathcal{J}_{\alpha}(\cdot ; \mathbf{m})\left(\boldsymbol{\Theta}_{\alpha}-\mathbf{m}_{\alpha}\right) \tag{A.28}
\end{equation*}
$$

and

$$
\begin{equation*}
\tilde{f}_{\beta}\left(\cdot ; \mathbf{m}, \boldsymbol{\Theta}_{\beta}\right) \doteq f(\cdot ; \mathbf{m})+\mathcal{J}_{\beta}(\cdot ; \mathbf{m})\left(\boldsymbol{\Theta}_{\beta}-\mathbf{m}_{\beta}\right) \tag{A.29}
\end{equation*}
$$

where $\mathcal{J}_{\alpha}(\cdot ; \mathbf{m})$ and $\mathcal{J}_{\beta}(\cdot ; \mathbf{m})$ are the columns of the Jacobian matrix corresponding to the sets of parameters $\alpha$ and $\beta$, respectively, and $\boldsymbol{\Theta}_{\alpha}$ and $\boldsymbol{\Theta}_{\beta}$ are the corresponding random parameter vectors.
Noting that Equation (A.27) expresses $\tilde{f}$ as a sum of (affine transformations of) random variables, we can use the fact that for independent Gaussian random variables $\mathbf{X}$ and $\mathbf{Y}$, the distribution $h_{\mathbf{Z}}$ of $\underset{\tilde{f}}{\mathbf{Z}}=\mathbf{X}+\mathbf{Y}$ is equal to the convolution of the distributions $h_{\mathbf{X}}$ and $h_{\mathbf{Y}}$ to obtain an approximation to $\tilde{f}$. In particular, for $\mathbf{Z}=\mathbf{X}+\mathbf{Y}$,

$$
\begin{equation*}
f_{\mathbf{Z}}(\mathbf{z})=\int_{-\infty}^{\infty} f_{\mathbf{Y}}(\mathbf{z}-\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) \mathrm{d} \mathbf{x} \tag{A.30}
\end{equation*}
$$

Letting $\mathbf{X}=\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right), \mathbf{Y}=\tilde{f}_{\beta}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\beta}\right)$, and $\mathbf{X}=\tilde{f}(\mathbf{X} ; \boldsymbol{\Theta})$, we can write

$$
\begin{align*}
& \tilde{g}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})) \\
& =\int_{-\infty}^{\infty} \tilde{g}_{\tilde{f}_{\beta}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\beta}\right)}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})-\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right)\right) \tilde{g}_{\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)}\left(\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right)\right) \mathrm{d} \mathbf{X}  \tag{A.31}\\
& =\int_{-\infty}^{\infty} \mathcal{N}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}) ; \boldsymbol{\mu}\left(\mathbf{X}, \mathbf{m}, \boldsymbol{\theta}_{\alpha}, \tilde{f}_{\alpha}\right), \boldsymbol{\Sigma}\left(\mathbf{X}, \mathbf{m}, \boldsymbol{\theta}_{\beta}, \mathbf{S}_{\beta}\right)\right) \tilde{g}_{\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)}\left(\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right)\right) \mathrm{d} \mathbf{X} \tag{A.32}
\end{align*}
$$

with

$$
\begin{equation*}
\boldsymbol{\mu}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}, \tilde{f}_{\alpha}\right)=f(\mathbf{X} ; \mathbf{m})+\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right) \tag{A.33}
\end{equation*}
$$

and

$$
\begin{equation*}
\boldsymbol{\Sigma}\left(\mathbf{X} ; \mathbf{m}, \mathbf{S}_{\beta}, \mathcal{J}_{\beta}\right)=\mathcal{J}_{\beta}(\mathbf{X} ; \mathbf{m}) \mathbf{S}_{\beta} \mathcal{J}_{\beta}(\mathbf{X} ; \mathbf{m})^{\top} \tag{A.34}
\end{equation*}
$$

where we have used the fact that for a Gaussian distribution with mean $m$ and covariance $S$, $\mathcal{N}\left(z-y ; m_{2} S\right)=\mathcal{N}(z ; m+y, S)$. We can then approximate the probability density function $\tilde{g}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}(\tilde{f}(\mathbf{X} ; \boldsymbol{\theta}))$ via the Monte Carlo estimator

$$
\begin{align*}
& \hat{\tilde{g}}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta})) \\
& =\frac{1}{R} \sum_{j=1}^{R} \mathcal{N}\left(\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}) ; \boldsymbol{\mu}\left(\mathbf{X}, \mathbf{m}, \tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right)^{(j)}\right), \mathbf{\Sigma}\left(\mathbf{X} ; \mathbf{m}, \mathbf{S}_{\beta}, \mathcal{J}_{\beta}\right)\right) \tag{A.35}
\end{align*}
$$

with $\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\theta}_{\alpha}\right)^{(j)} \sim \tilde{g}_{\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)}$. Finally, we can express the distribution $\hat{\tilde{g}}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}$ as

$$
\begin{equation*}
\hat{\tilde{g}}_{\tilde{f}(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta})}=\frac{1}{R} \sum_{j=1}^{R} \mathcal{N}\left(f(\mathbf{X} ; \mathbf{m})+\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)^{(j)}, \mathcal{J}_{\beta}(\mathbf{X} ; \mathbf{m}) \mathbf{S}_{\beta} \mathcal{J}_{\beta}(\mathbf{X} ; \mathbf{m})^{\top}\right) \tag{A.36}
\end{equation*}
$$

where $g_{\boldsymbol{\Theta}_{\beta}}=\mathcal{N}\left(\mathbf{m}_{\beta}, \mathbf{S}_{\beta}\right)$ and samples $\tilde{f}_{\alpha}\left(\mathbf{X} ; \mathbf{m}, \boldsymbol{\Theta}_{\alpha}\right)^{(j)}$ are obtained by sampling parameters from the distribution $g_{\boldsymbol{\Theta}_{\alpha}}=\mathcal{N}\left(\mathbf{m}_{\alpha}, \mathbf{S}_{\alpha}\right)$. This concludes the proof.

## Appendix B Further Empirical Results

## B. 1 Tabular Results for Diabetic Retinopathy Diagnosis Tasks

The results below were reproduced from Band et al. [2021] using the RETINA benchmark.
Table 3: Country Shift. Prediction and uncertainty quality of baseline methods in terms of the area under the receiver operating characteristic curve (AUC) and classification accuracy, as a function of the proportion of data referred to a medical expert. All methods are tuned on in-domain validation AUC, and ensembles have $K=3$ constituent models (true for all subsequent tables unless specified otherwise). On in-domain data, MC DROPOUT performs best across all thresholds. On distributionally shifted data, no method consistently performs best.

| Method | No Referral |  | 50\% Data Referred |  | 70\% Data Referred |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | AUC (\%) $\uparrow$ | Accuracy (\%) $\uparrow$ | AUC (\%) $\uparrow$ | Accuracy (\%) $\uparrow$ | AUC (\%) $\uparrow$ | Accuracy $\uparrow$ |
| EyePACS Dataset (In-Domain) |  |  |  |  |  |  |
| MAP (Deterministic) | $87.4 \pm 1.3$ | $88.6 \pm 0.7$ | $91.1 \pm 1.8$ | $95.9 \pm 0.4$ | $94.9 \pm 1.1$ | $96.5 \pm 0.3$ |
| MFVI | $83.3 \pm 0.2$ | $85.7 \pm 0.1$ | $85.5 \pm 0.7$ | $94.5 \pm 0.1$ | $88.2 \pm 0.7$ | $95.9 \pm 0.1$ |
| RADIAL-MFVI | $83.2 \pm 0.5$ | $74.2 \pm 5.0$ | $88.9 \pm 0.9$ | $81.8 \pm 6.0$ | $91.2 \pm 1.3$ | $83.8 \pm 5.5$ |
| FSVI | $88.5 \pm 0.1$ | $89.8 \pm 0.0$ | $91.0 \pm 0.4$ | $96.4 \pm 0.0$ | $94.3 \pm 0.3$ | $97.2 \pm 0.1$ |
| MC DROPOUT | $91.4 \pm 0.2$ | $90.9 \pm 0.1$ | $95.3 \pm 0.2$ | $97.4 \pm 0.1$ | $97.4 \pm 0.1$ | $98.1 \pm 0.0$ |
| RANK-1 | $85.6 \pm 1.4$ | $87.7 \pm 0.8$ | $87.1 \pm 2.3$ | $95.3 \pm 0.5$ | $90.9 \pm 2.0$ | $96.4 \pm 0.4$ |
| DEEP ENSEMBLE | $90.3 \pm 0.2$ | $90.3 \pm 0.3$ | $91.7 \pm 0.6$ | $97.2 \pm 0.0$ | $95.0 \pm 0.5$ | $97.9 \pm 0.0$ |
| MFVI ENSEMBLE | $85.4 \pm 0.0$ | $87.8 \pm 0.0$ | $86.3 \pm 0.4$ | $95.4 \pm 0.0$ | $89.2 \pm 0.4$ | $96.7 \pm 0.1$ |
| RADIAL-MFVI ENSEMBLE | $84.9 \pm 0.1$ | $74.2 \pm 1.5$ | $91.4 \pm 0.2$ | $83.4 \pm 1.7$ | $93.3 \pm 0.3$ | $85.9 \pm 1.6$ |
| FSVI ENSEMBLE | $90.3 \pm 0.1$ | $90.6 \pm 0.0$ | $92.1 \pm 0.2$ | $97.1 \pm 0.0$ | $95.2 \pm 0.2$ | $97.8 \pm 0.1$ |
| MC DROPOUT ENSEMBLE | $\mathbf{9 2 . 5} \pm \mathbf{0 . 0}$ | $\mathbf{9 1 . 6} \pm \mathbf{0 . 0}$ | $\mathbf{9 5 . 8} \pm \mathbf{0 . 1}$ | $\mathbf{9 7 . 8} \pm \mathbf{0 . 0}$ | $\mathbf{9 7 . 7} \pm \mathbf{0 . 1}$ | $\mathbf{9 8 . 4} \pm \mathbf{0 . 0}$ |
| RANK-1 ENSEMBLE | $89.5 \pm 0.8$ | $89.3 \pm 0.4$ | $88.5 \pm 1.3$ | $96.9 \pm 0.3$ | $91.6 \pm 1.2$ | $97.6 \pm 0.3$ |
| APTOS 2019 Dataset (Population Shift) |  |  |  |  |  |  |
| MAP (Deterministic) | $92.2 \pm 0.2$ | $86.2 \pm 0.6$ | $80.1 \pm 3.6$ | $87.6 \pm 1.5$ | $55.4 \pm 4.3$ | $85.4 \pm 1.2$ |
| MFVI | $91.4 \pm 0.2$ | $84.1 \pm 0.3$ | $93.8 \pm 0.4$ | $92.1 \pm 0.5$ | $93.0 \pm 0.6$ | $92.7 \pm 0.5$ |
| RADIAL-MFVI | $90.7 \pm 0.7$ | $71.8 \pm 4.6$ | $82.0 \pm 2.5$ | $81.5 \pm 2.7$ | $66.4 \pm 2.1$ | $85.9 \pm 1.0$ |
| FSVI | $94.1 \pm 0.1$ | $87.6 \pm 0.5$ | $90.6 \pm 0.9$ | $90.7 \pm 0.7$ | $77.2 \pm 4.6$ | $89.8 \pm 0.3$ |
| MC DROPOUT | $94.0 \pm 0.2$ | $86.8 \pm 0.2$ | $87.4 \pm 0.3$ | $88.1 \pm 0.2$ | $65.3 \pm 1.7$ | $88.2 \pm 0.4$ |
| RANK-1 | $92.5 \pm 0.3$ | $86.2 \pm 0.5$ | $90.1 \pm 2.5$ | $91.4 \pm 1.1$ | $75.1 \pm 7.8$ | $89.5 \pm 1.5$ |
| DEEP ENSEMBLE | $94.2 \pm 0.2$ | $87.5 \pm 0.1$ | $91.2 \pm 1.9$ | $92.4 \pm 0.9$ | $67.4 \pm 7.3$ | $90.1 \pm 1.2$ |
| MFVI ENSEMBLE | $93.2 \pm 0.1$ | $87.0 \pm 0.2$ | $\mathbf{9 4 . 9} \pm \mathbf{0 . 3}$ | $\mathbf{9 3 . 7} \pm 0.3$ | $\mathbf{9 4 . 2} \pm \mathbf{0 . 3}$ | $\mathbf{9 4 . 0} \pm 0.3$ |
| RADIAL-MFVI ENSEMBLE | $91.8 \pm 0.2$ | $69.0 \pm 1.9$ | $78.6 \pm 0.6$ | $79.8 \pm 0.9$ | $60.9 \pm 0.3$ | $86.7 \pm 0.2$ |
| FSVI ENSEMBLE | $\mathbf{9 4 . 6} \pm 0.1$ | $88.9 \pm 0.2$ | $90.7 \pm 0.5$ | $91.1 \pm 0.6$ | $74.1 \pm 3.4$ | $89.8 \pm 0.2$ |
| MC DROPOUT ENSEMBLE | $94.1 \pm 0.1$ | $87.6 \pm 0.1$ | $86.8 \pm 0.2$ | $88.0 \pm 0.2$ | $62.3 \pm 0.4$ | $87.7 \pm 0.2$ |
| RANK-1 ENSEMBLE | $94.1 \pm 0.2$ | $88.3 \pm 0.2$ | $\mathbf{9 4 . 9} \pm \mathbf{0 . 4}$ | $93.5 \pm 0.3$ | $92.4 \pm 1.5$ | $93.8 \pm 0.3$ |

## B. 2 UCI Regression

Table 4: This table compares the predictive performance between the method proposed in this paper and the method proposed by Sun et al. [2019] on six datasets from the UCI database. We followed the same training protocol as Sun et al. [2019] and used the code provided by the authors to load and process the data. The same network architecture was used (one hidden layer with 50 hidden units). We report the results for the best set of hyperparameters, computed over ten random seeds. Lower RMSE and higher log-likelihood are better. Best results are shaded in gray. The first five rows are small-scale UCI experiments, and the sixth row ("Protein") is a larger-scale experiment ( 45,740 data points).

|  | RMSE |  | Log-Likelihood |  |
| :--- | :---: | :---: | :---: | :---: |
|  | Sun et al. [2019] | Ours | Sun et al. [2019] | Ours |
| Boston | $\mathbf{2 . 3 7 8} \pm \mathbf{0 . 1 0 4}$ | $3.632 \pm 0.515$ | $-\mathbf{2 . 3 0 1} \pm \mathbf{0 . 0 3 8}$ | $-3.150 \pm 0.495$ |
| Concrete | $4.935 \pm 0.180$ | $\mathbf{4 . 1 7 7} \pm \mathbf{0 . 4 4 3}$ | $-3.096 \pm 0.016$ | $-\mathbf{2 . 8 5 5} \pm \mathbf{0 . 1 1 6}$ |
| Energy | $0.412 \pm 0.017$ | $\mathbf{0 . 4 0 9} \pm \mathbf{0 . 0 6 0}$ | $-0.684 \pm 0.020$ | $-\mathbf{0 . 5 3 9} \pm \mathbf{0 . 1 3 8}$ |
| Wine | $0.673 \pm 0.014$ | $\mathbf{0 . 6 1 5} \pm \mathbf{0 . 0 3 3}$ | $-1.040 \pm 0.013$ | $-\mathbf{0 . 9 5 9} \pm \mathbf{0 . 0 3 4}$ |
| Yacht | $0.607 \pm 0.068$ | $\mathbf{0 . 5 1 4} \pm \mathbf{0 . 2 4 2}$ | $-1.033 \pm 0.033$ | $-\mathbf{0 . 8 8 8} \pm \mathbf{0 . 3 3 4}$ |
| Protein | $4.326 \pm 0.019$ | $\mathbf{4 . 2 4 8} \pm \mathbf{0 . 0 4 3}$ | $-2.892 \pm 0.004$ | $\mathbf{- 2 . 8 6 6} \pm \mathbf{0 . 0 0 9}$ |

## Appendix C Illustrative Examples

## C. 1 Two Moons Classification Task



Figure 5: Binary classification on the Two Moons dataset. The plots show the posterior predictive mean and variance of a BNN trained via FSVI (Figure 5a and Figure 5b), of a BNN trained via MFVI (Figure 5c and Figure 5d), and an ensemble of MAP models (Figure 5e and Figure 5f). The predictive means represent the expected class probabilities and the predictive variance the model's epistemic uncertainty over the class probabilities. With FSVI, the predictive distribution is able to faithfully capture the geometry of the data manifold and exhibits high uncertainty over the class probabilities in areas of the data space of which the data is not informative. In contrast, neither MFVI, nor MAP ensembles are unable to accurately capture the geometry of the data manifold only exhibit high uncertainty around the decision boundary.

## C. 2 Synthetic 1D Regression Datasets



Figure 6: 1D Regression with FSVI on a selection of datasets used to demonstrate desirable predictive uncertainty estimates in prior works. The left column is zoomed in.

## Appendix D Implementation, Training, and Evaluation Details

## D. 1 Hyperparameter Selection Protocol

For FSVI, we used a holdout validation set ( $10 \%$ of the training set) to conduct a hyperparameter search over the prior variance, the number of context points used to evaluate the KL divergence, the context distribution, and the number of Monte Carlo samples used to evaluate the expected loglikelihood. We selected the set of hyperparameters that yielded the highest validation log-likelihood for all experiments. We state the hyperparameters selected for the different datasets below.
For other methods, we used a holdout validation set of the same size and selected the best-performing hyperparameters. We used implementations provided by the authors of MFVI (radial) and SWAG. All other methods were implemented from scratch unless stated otherwise.

## D. 2 FashionMNIST vs. MNIST/NotMNIST

We train all model on the FashionMNIST dataset and evaluate the models' predictive uncertainty performance on out-of-distribution data on the MNIST dataset. Both datasets consist of images of size $28 \times 28$ pixels. The FashionMNIST dataset is normalized to have zero mean and a standard deviation of one. The MNIST dataset is normalized with the same transformation, that is, using the same mean and standard deviation used for the in-distribution data. We chose FashionMNIST/MNIST instead of MNIST/NotMNIST because the latter is notably easier than the former.
In this experiment, a network architecture with two convolutional layers of 32 and $643 \times 3$ filters and a fully-connected final layer of 128 hidden units is used. A max pooling operation is placed after each convolutional layer and ReLU activations are used. We do not use batch normalization. All models are trained for 30 epochs with a mini-batch size of 128 using SGD with a learning rate of $5 \times 10^{-3}$, momentum (with momentum parameter 0.9 ), and a cosine learning rate schedule with parameter 0.05.
For FSVI with $p_{\mathbf{X}_{\mathcal{C}}}=$ random monochrome, we sampled $50 \%$ of the context points for each gradient step from the mini-batch and the other $50 \%$ according to the method described in Appendix D.8. For FSVI with $p_{\mathbf{X}_{\mathcal{C}}}=$ KNIST, we used the KMNIST dataset.

## D. 3 CIFAR-10 vs. SVHN

We train all model on the CIFAR-10 dataset and evaluate the models' predictive uncertainty performance on out-of-distribution data on the SVHN dataset. Both datasets consist of images of size $32 \times 32 \times 3$, with RBG channels. The CIFAR-10 dataset is normalized to have zero mean and a standard deviation of one. The SVHN dataset is normalized with the same transformation, that is, using the same mean and standard deviation used for the in-distribution data. The training data is augmented with random horizontal flips (with a probability of 0.5 ) and random crops ( 4 zero pixels on all sides).
In this experiment, a standard ResNet-18 network architecture was used. All models are trained for 200 epochs with a mini-batch size of 128 using SGD with a learning rate of $5 \times 10^{-3}$, momentum (with momentum parameter 0.9), and a cosine learning rate schedule with parameter 0.05.
For FSVI with $p_{\mathbf{X}_{\mathcal{C}}}=$ random monochrome, we sampled $100 \%$ of the context points for each gradient step from the mini-batch and the other $50 \%$ according to the method described in Appendix D.8. For FSVI with $p_{\mathbf{X}_{\mathcal{C}}}=$ CIFAR-100, we used the CIFAR-100 dataset.

## D. 4 Diabetic Retinopathy Diagnosis

Prediction and Expert Referral. In real-world settings where the evaluation data may be sampled from a shifted distribution, incorrect predictions may become increasingly likely. To account for that possibility, predictive uncertainty estimates can be used to identify datapoints where the likelihood of an incorrect prediction is particularly high and refer them for further review. We consider a corresponding selective prediction task, where the predictive performance of a given model is evaluated for varying expert referral rates. That is, for a given referral rate of $\gamma \in[0,1]$, a model's predictive uncertainty is used to identify the $\gamma$ proportion of images in the evaluation set for which the model's predictions are most uncertain. Those images are referred to a medical professional for further
review, and the model is assessed on its predictions on the remaining $(1-\gamma)$ proportion of images. By repeating this process for all possible referral rates and assessing the model's predictive performance on the retained images, we estimate how reliable it would be in a safety-critical downstream task, where predictive uncertainty estimates are used in conjunction with human expertise to avoid harmful predictions. Importantly, selective prediction tolerates out-of-distribution examples. For example, even if unfamiliar features appear in certain images, a model with reliable uncertainty estimates will perform better in selective prediction by assigning these images high epistemic (and predictive) uncertainty, therefore referring them to an expert at a lower $\gamma$.

For all methods, experiments are performed using a ResNet-50 network architecture. Training and evaluation scripts as well as model checkpoints can be found at

```
github.com/google/uncertainty-baselines/.../diabetic_retinopathy_detection.
```


## D. 5 Two Moons

In this experiment, we use a multi-layer perceptron (MLP) consisting of two fully-connected layers with 30 hidden units each and tanh activations. We train all models with a learning rate of $10^{-3}$.

For FSVI, we sampled context points uniformly from $[-10,10] \times[-10,10]$.

## D. 6 1D Regression

In this experiment, we use a multi-layer perceptron (MLP) consisting of two fully-connected layers with 100 hidden units each and ReLU activations.

For FSVI, we sampled context points uniformly from $[-10,10]$.

## D. 7 Further Implementation Details

We use the Adam optimizer with default settings of $\beta_{1}=0.9, \beta_{2}=0.99$ and $\epsilon=10^{-8}$ for all experiments. The deterministic neural networks that were used for the ensemble were trained with a weight decay of $\lambda=1 \mathrm{e}-1$. MFVI (tempered) was trained with a KL scaling factor of 0.1 to obtain a cold posterior.

## D. 8 Selection of Context Distribution

We estimate the supremum at every gradient step by sampling a set of context points $\mathbf{X}_{\mathcal{C}}$ from a distribution $p_{\mathbf{X}_{\mathcal{C}}}$ at every gradient step. For tasks with image inputs, we construct a distribution $p_{\mathbf{X}_{\mathcal{C}}}$, defined as a uniform distribution over images with monochromatic channels. To generate a sample from this "monochrome images" distribution, we first take all images in the training data, flatten each channel, and stack the flattened image channels into a single vector each. We then draw a random element (i.e., a pixel) from each channel vector and then use these pixels to generate a monochrome image of a given resolution by setting every channel equal to the value of the pixel that was drawn. For regression tasks with a $D$-dimensional input space, $p_{\mathbf{X}_{\mathcal{C}}}$ is defined as a uniform distribution with lower and upper bounds set to the empirical lower and upper bounds of the training data. For further details on the effect of different sampling schemes on the posterior predictive distribution's performance, see Appendix B.

## D. 9 Compute Resources

All experiments were carried out on an Nvidia V-100 GPU with 32GB of memory.

