
Last Layer Marginal Likelihood for Invariance Learning

Pola Schwöbel

Technical University of Denmark

Martin Jørgensen

University of Oxford

Sebastian W. Ober

University of Cambridge

Mark van der Wilk

Imperial College London

Abstract

Data augmentation is often used to incorporate inductive biases into models. Traditionally, these are hand-crafted and tuned with cross validation. The Bayesian paradigm for model selection provides a path towards end-to-end learning of invariances using only the training data, by optimising the marginal likelihood. Computing the marginal likelihood is hard for neural networks, but success with tractable approaches that compute the marginal likelihood for the last layer only raises the question of whether this convenient approach might be employed for learning invariances. We show partial success on standard benchmarks, in the low-data regime and on a medical imaging dataset by designing a custom optimisation routine. Introducing a new lower bound to the marginal likelihood allows us to perform inference for a larger class of likelihood functions than before. On the other hand, we demonstrate failure modes on the CIFAR10 dataset, where the last layer approximation is not sufficient due to the increased complexity of our neural network. Our results indicate that once more sophisticated approximations become available the marginal likelihood is a promising approach for invariance learning in neural networks.

1 INTRODUCTION

Human learners generalise from example to category with seemingly little effort. Machine learning models

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

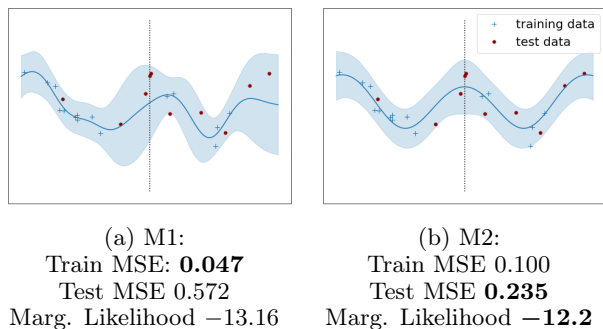


Figure 1: A non-invariant model M1 and its sign invariant (i.e. symmetric around $x = 0$) counterpart M2. The non-invariant M1 has a better train MSE, but the invariant M2 has a better test MSE. The log marginal likelihood correctly identifies M2 as better.

aim to make accurate predictions on unseen data points based on finitely many examples. This generalisation is enabled by *inductive biases*. In *Steps toward Artificial Intelligence* Marvin Minsky (1961) highlights the importance of invariance as an inductive bias: ‘*One of the prime requirements of a good property is that it be invariant under the commonly encountered equivalence transformations. Thus for visual Pattern-Recognition we would usually want the object identification to be independent of uniform changes in size and position.*’ In modern machine learning pipelines invariances are achieved through data augmentation. If we, for example, would like our neural network to be invariant with respect to rotation, we simply present it with rotated versions of the input data. Data augmentation schemes are almost always hand-crafted, based on assumptions and expert knowledge about the data, or found by cross-validation. We aim to learn invariances with backpropagation, to reduce the human intervention in the design of ML algorithms.

Learning invariances through gradients requires a suit-

able loss function. Standard losses like negative log-likelihood or mean squared error solely measure how tightly we fit the training data. Good inductive biases (e.g. convolutions) constrain the expressiveness of a model, and therefore do not improve the fit on the training data. Thus, they can not be learned by minimising the training loss alone.

In Bayesian inference, this problem is known as *model selection*, and is commonly solved by using a different training objective: the marginal likelihood. For a model of data y , parametrised by weights w and hyperparameters θ it is given by

$$p(y|\theta) = \int p(y|w)p(w|\theta)dw. \quad (1)$$

As opposed to standard training losses, it correlates with generalisation, and thus provides a general way to select an inductive bias, independent of parameterisation (Williams and Rasmussen, 2006; Rasmussen and Ghahramani, 2001; MacKay, 2003). Van der Wilk et al. (2018) demonstrated that invariances can be learned by straightforward backpropagation using the marginal likelihood in Gaussian process (GP) models, where the marginal likelihood can be accurately approximated. Fig. 1 shows an invariant and a non-invariant GP; the invariant model has higher marginal likelihood as well as lower test mean squared error. Thus, the marginal likelihood correctly identifies invariance as a useful inductive bias.

Current GP models often lack predictive performance compared to their highly expressive neural network counterparts, hence applying this elegant principle to neural networks is attractive. The challenge is, however, that finding accurate and differentiable marginal likelihood approximations for neural networks is still an open problem. In this work we investigate a convenient short-cut: computing Bayesian quantities only in the last layer. This avoids difficulties of the marginal likelihood in the full network, and has already been shown helpful (Wilson et al., 2016a,b). Given the possible impact of invariance learning with the convenience of the last-layer approximation, it is important to investigate its potential. Our results provide a nuanced picture of this approach: there are situations where the last-layer approximation is sufficient, but others where it is not.

To provide these results, we

1. construct a *deep neural network with a Bayesian last layer that incorporates invariance*, based on invariant GPs (van der Wilk et al., 2018) and deep kernel learning (Wilson et al., 2016b),
2. overcome problems with the training implied by a straightforward combination of Van der Wilk

et al. (2018) and Wilson et al. (2016b) via a *new optimisation scheme*, and a *new variational bound* that allows for non-Gaussian likelihoods,

3. *investigate failure modes* on more complex model architectures to show limitations of using the last-layer approximation for invariance learning.

2 RELATED WORK

Bayesian Deep Learning aims to provide principled uncertainty quantification for deep models. Exact computation for Bayesian deep models is intractable, so different approximations have been suggested. Variational strategies (e.g. Blundell et al., 2015) maximise the evidence lower bound (ELBO) to the marginal likelihood, thereby minimising the gap between approximate and true posteriors. To remain computationally feasible, approximations for Bayesian neural networks are often crude, and while weight posteriors are useful in practice, the marginal likelihood estimates are typically too imprecise for hyperparameter estimation (Blundell et al., 2015; Turner and Sahani, 2011). Hyperparameter estimation in deep GPs has achieved more success (Damianou and Lawrence, 2013; Dutordoir et al., 2020), but training deep GPs can be challenging. Some very recent works have shown initial promise in using the marginal likelihood for hyperparameter selection in Bayesian neural networks (Ober and Aitchison, 2020; Immer et al., 2021; Dutordoir et al., 2021). Instead of a Bayesian treatment of all weights using rough approximations, we follow a deep kernel learning approach, i.e. computing the marginal likelihood for the last layer only.

Deep Kernel Learning (DKL; Hinton and Salakhutdinov, 2007; Calandra et al., 2016; Bradshaw et al., 2017) replaces the last layer of a neural network with a GP, where marginal likelihood estimation is accurate (Burt et al., 2020). Wilson et al. (2016a,b) had significant success achieving improved uncertainty estimates. Their results indicate that such a neural network-GP hybrid is promising for invariance learning. Ober et al. (2021) identify difficulties with overfitting in DKL models, but also show mechanisms by which such overfitting is mitigated. We find similar issues and adapt the standard DKL training procedure to avoid them when learning invariance hyperparameters. We will discuss these issues in more depth as we describe our training procedure in Sec. 5.

Data Augmentation is used to incorporate invariances into deep learning models. Where good invariance assumptions are available a priori (e.g. for natural images) this improves generalisation performance and is ubiquitous in deep learning pipelines. Instead of relying on assumptions and hand-crafting, recent approaches

learn data augmentation schemes. Cubuk et al. (2019, 2020) and Ho et al. (2019) train on the validation data, and use reinforcement learning and evolutionary search respectively to find parameters. Zhou et al. (2021); Lorraine et al. (2020) compute losses on validation sets for learning invariance parameters, and estimate gradients w.r.t. them in outer loops. Similar to our work, Benton et al. (2020) learn data augmentations on training data end-to-end, by adding a regularisation term to the negative log-likelihood loss that encourages invariance. They argue that tuning this regularisation term via cross-validation can be avoided, since the loss function is relatively flat. Yet, the method relies on explicit regularisation, and thus on an understanding of the parameters in question. Our method is based on a Bayesian view of data augmentation as incorporating an invariance on the functions in the prior distribution (van der Wilk et al., 2018; Nabarro et al., 2021). This allows the marginal likelihood to be used as an objective for learning invariances. This has many advantages, such as allowing backpropagation from training data, automatic and principled regularisation, and parameterisation independence (see Sec. 5). This makes the marginal likelihood objective a promising avenue for future work, which may want to incorporate invariances whose parameterisations are non-interpretable.

3 BACKGROUND

3.1 Variational Gaussian processes

A Gaussian process (GP) (Williams and Rasmussen, 2006) is a distribution on functions with the property that any vector of function values $\mathbf{f} = (f(x_1), \dots, f(x_N))$ is Gaussian distributed. We assume zero mean functions and real valued vector inputs.

Inference in GP models with general likelihoods and big datasets can be done with variational approximations (Titsias, 2009; Hensman et al., 2015). The approximate posterior is constructed by conditioning the prior on M inducing variables $\mathbf{u} \in \mathbb{R}^M$, and specifying their marginal distribution with $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ (for overviews see Bui et al. 2017; van der Wilk et al. 2020). This results in a variational predictive distribution:

$$q(f(x^*)) = \mathcal{N}(\boldsymbol{\alpha}(x^*)^\top \mathbf{m}, k(x^*, x^*) - \boldsymbol{\alpha}(x^*)^\top (\mathbf{K}_{zz} - \mathbf{S}) \boldsymbol{\alpha}(x^*)), \quad (2)$$

where $\mathbf{z} \in \mathbb{R}^{M \times d}$ are inducing inputs, \mathbf{K}_{zz} is the matrix with entries $k(z_i, z_j)$, $\boldsymbol{\alpha}(x^*) = \mathbf{K}_{zz}^{-1} k(\mathbf{z}, x^*)$, and k is the chosen covariance function.

Variational inference (VI) selects an approximation by minimising the KL divergence of the approximation to the true posterior with respect to the variational parameters $\mathbf{z}, \mathbf{m}, \mathbf{S}$. This is done by maximising a

lower bound to the marginal likelihood (the “evidence”), which has the KL divergence as its gap (Matthews et al., 2016). The resulting evidence lower bound (ELBO) is

$$\log p(y) \geq \mathcal{L} = \sum_{n=1}^N \mathbb{E}_{q(f(x_n))} [\log p(y_n | f(x_n))] - \text{KL}[q(\mathbf{u}) || p(\mathbf{u})]. \quad (3)$$

In exact GPs, (kernel) hyperparameters are found by maximising the log marginal likelihood $\log p(y)$ (Williams and Rasmussen, 2006). For our models of interest, the exact marginal likelihood is intractable. We use the ELBO as a surrogate. This results in an approximate inference procedure that maximises the ELBO with respect to both the variational parameters and the hyperparameters. Optimising the variational parameters improves the quality of the posterior approximation, and tightens the bound to the marginal likelihood. Optimising the hyperparameters hopefully improves the model, but the slack in the ELBO can lead to worse hyperparameter selection (Turner and Sahani, 2011).

3.2 Invariant Gaussian Processes

A function $f : \mathcal{X} \rightarrow \mathcal{Y}$ is *invariant* to a transformation $t : \mathcal{X} \rightarrow \mathcal{X}$ if $f(x) = f(t(x))$, $\forall x \in \mathcal{X}$, and $\forall t \in \mathcal{T}$. I.e., an invariant function will have the same output for a certain range of transformed inputs known as the *orbit*. A straightforward way to construct invariant functions is to simply average a function over the orbit (Kondor, 2008; Ginsbourger et al., 2012, 2013). We consider a similar construction where we average a function over a data augmentation distribution, which results in an approximately invariant function where $f(x) \approx f(t(x))$ (van der Wilk et al., 2018; Dao et al., 2019). Augmented data samples x_a are obtained by applying random transformations t to an input, $x_a = t(x)$, leading to the distribution $p(x_a|x)$. That is, an approximately invariant function f can be constructed from any non-invariant g as

$$f(x) = \sum_{t \in \mathcal{T}} g(t(x)), \text{ or } f(x) = \int g(x_a) p(x_a|x) dx_a. \quad (4)$$

Van der Wilk et al. (2018) exploit this construction to build a GP with continuously adjustable invariances. They place a GP prior on $g \sim \mathcal{GP}(0, k_g(\cdot, \cdot))$, and since Gaussians are closed under summations, f is a GP too. By construction f is invariant to the augmentation distribution $p(x_a|\cdot)$ and its kernel is given by

$$k_f(x, x') = \iint k_g(x_a, x'_a) p(x_a|x) p(x'_a|x') dx_a dx'_a. \quad (5)$$

Non-trivial $p(x_a|x)$ densities present a problem for standard VI, as the kernel evaluations in eq. 2 become

intractable. This is solved by making the inducing variables observations of g rather than the usual f . This ensures that $\mathbf{K}_{\mathbf{z}\mathbf{z}}$ is tractable, as it only requires evaluations of k_g , which makes the KL divergence tractable. When the likelihood is Gaussian, it additionally provides a way to tackle the expected log likelihood:

$$\mathbb{E}_{g(f(x))} \log \mathcal{N}(y; f(x), \sigma^2) = \text{const} - \frac{(y_n - \mu)^2 + \tau}{2\sigma^2} \quad (6)$$

where μ, τ are the mean and variance in (2). Only unbiased estimates of μ, μ^2 and τ are needed for an unbiased estimate of the ELBO. These can be obtained from simple Monte Carlo estimates of k_f (5), and $k(\mathbf{z}, x)$.¹

3.3 Parameterising learnable invariances

The invariance of the GP in (5) is learned by adjusting the augmentation distribution. We parameterise the distribution and treat its parameters as kernel hyperparameters. We learn these by maximising the ELBO. As done in similar work (Benton et al., 2020; van der Wilk et al., 2018), we consider affine transformations. Our affine transformations are controlled by $\phi = (\alpha, s^x, s^y, p^x, p^y, t^x, t^y)$, which describes rotation, scale, shearing and horizontal and vertical translation. We parameterise a family of augmentation distributions by specifying uniform ranges with $\phi_{\min}, \phi_{\max} \in \mathbb{R}^7$ that are to be applied to the input image. Different ranges that are learned on ϕ_{\min}, ϕ_{\max} correspond to different invariances in $f(\cdot)$. For example, learning $\phi_{\min/\max} = (\pm\pi, 0, 0, 0, 0, 0, 0)$ corresponds to full rotational invariance (sampling any angle between $-\pi$ and π) but no scaling, shearing or translations.

We sample from the resulting $p(x_a|x, \phi_{\max}, \phi_{\min})$ (we will write $p(x_a|x, \phi)$ for brevity) by **1**) sampling the parameters for a transformation from a uniform distribution, **2**) generating a transformed coordinate grid, and **3**) interpolating² the image x :

$$x_a = t_\nu(x), \quad \nu \sim U(-\phi_{\min}, \phi_{\max}). \quad (7)$$

Since transforming $t_\nu(x)$ is differentiable, this procedure is reparameterisable w.r.t. ϕ_{\max}, ϕ_{\min} via $\nu = \phi_{\min} + (\phi_{\max} - \phi_{\min})\varepsilon$, $\varepsilon \sim U(0, 1)$. Straightforward automatic differentiation of the unbiased ELBO estimator described in the previous section provides the required gradients.

In summary, we learn $\phi_{\min/\max}$ by maximising the ELBO, so the transformations and their magnitudes are learned based on the specific training set. Different invariances will be learned for different training

¹We obtain $k(\mathbf{z}, x) = \int k_g(\mathbf{z}, x_a) p(x_a|x) dx_a$ from the interdomain trick, which can be estimated with Monte Carlo. See Van der Wilk et al. (2018) for details.

²Image transformation code from [gi thub.com/kevinzakka/spatial-transformer-network](https://github.com/kevinzakka/spatial-transformer-network)

data. The next sections show how these principles have potential even in neural network models, beyond the single layer GPs of Van der Wilk et al. (2018).

Algorithm 1: InvDKGP forward pass

1. Draw S samples from the augmentation distribution $x_a^i \sim p(x_a|x, \phi)$, $i = 1 \dots S$.
 2. Pass the x_a^i through the neural net h_w .
 3. Map extracted features using the non-inv. g .
 4. Aggregate samples to obtain inv. $f(x)$ by
 - (i) using the unbiased estimators from Sec. 3.2 in the Gaussian case, or,
 - (ii) averaging predictions $g(h_w(x_a^i))$, $i = 1, \dots, S$ directly in the Softmax case, see (16).
-

4 MODEL

As discussed in Sec. 1, we aim to learn neural network (NN) invariances through backpropagation, in the same way as is possible for single-layer GPs. Since finding high-quality approximations to the marginal likelihood of a NN is an ongoing research problem, we investigate whether a simpler *deep kernel* approach is sufficient. This uses a GP as the last layer of a NN, and takes advantage of accurate marginal likelihood approximations for the GP last layer. Success with such a simple method would significantly help automatic adaptation of data augmentation in neural network models. We hypothesise that the last layer approximation is sufficient, since data augmentation influences predictions only in the last layer (in the sense that one can construct an invariant function f from an arbitrary non-invariant g by summing in the last layer, eq. 4). See Fig. 2 for a graphical representation and Algorithm 1 for forward pass computations.

Deep Kernels take advantage of covariance functions being closed under transformations of their input. That is, if $k_g(\cdot, \cdot)$ is a covariance function on $\mathbb{R}^D \times \mathbb{R}^D$, then $k_g(h_w(\cdot), h_w(\cdot))$ is a covariance function on $\mathbb{R}^d \times \mathbb{R}^d$ for mappings $h_w : \mathbb{R}^d \rightarrow \mathbb{R}^D$. In our case, h_w is a NN parametrised by weights w , and hence w are viewed as hyperparameters of the kernel. The GP prior becomes

$$p(g) = \mathcal{GP}(0, k_g(h_w(\cdot), h_w(\cdot))). \quad (8)$$

The idea is to learn w along with the kernel hyperparameters. Importantly, this model remains a GP and so the inference described in Sec. 3 applies.

Our invariant model combines the flexibility of a NN $h_w(\cdot)$ with a GP g in the last layer, while ensuring

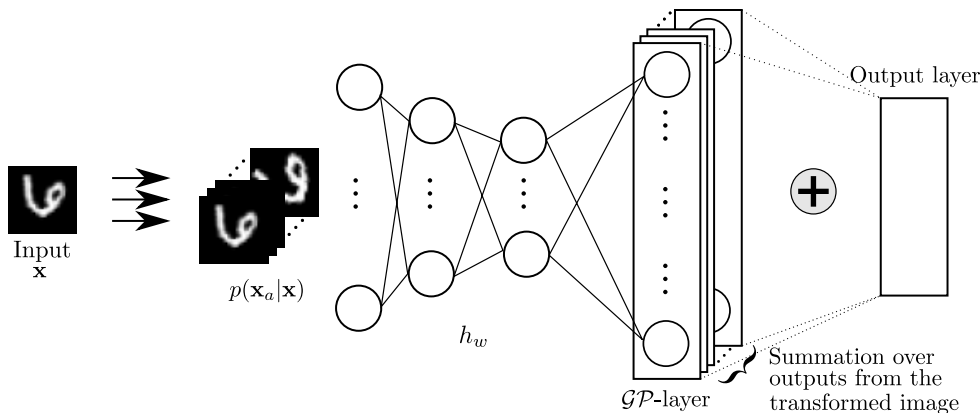


Figure 2: A visualisation of the model pipeline. For any input x , we can sample from the orbit distribution $p(x_a|x, \phi)$; each of these sample gets passed through a neural network parametrised by w . The last layer is a of the net is a GP, on which we can sum across sample outputs to create an invariant function.

overall invariance using the construction from (4):

$$f(x) = \int g(h_w(x_a))p(x_a|x, \phi)dx_a. \quad (9)$$

Thus, combining (5) and (9), f is an *invariant* GP with a *deep* kernel given as

$$k_f(x, x') = \int k_g(h_w(x_a), h_w(x'_a)) p(x_a|x, \phi)p(x'_a|x', \phi)dx_adx'_a. \quad (10)$$

The model is trained to fit observations y through the likelihood function $p(y|f(x))$, where we assume observations y_i are independent conditioned on the marginals $f(x_i)$.

Initially, we investigate training a model by simply combining the invariant GP training objective for Gaussian likelihoods (van der Wilk et al., 2018) with standard deep kernel learning (Wilson et al., 2016a,b). However, as we will discuss, several issues prevent these training procedures from working. In following sections we investigate why, provide solutions, and introduce a new ELBO that is suitable for more general likelihoods which improves training behaviour. We refer to our model as the *Invariant Deep Kernel GP (InvDKGP)*. An implementation can be found at <https://github.com/polaschwoebel/InvDKGP>.

5 DESIGNING A TRAINING SCHEME

The promise of deep kernel learning as presented by Wilson et al. (2016a,b) lies in training the NN and GP hyperparameters *jointly*, using the marginal likelihood as for standard GPs.³ However, prior works have noted

³Given that this quantity is difficult to approximate, we verify experimentally that we indeed need it and cannot

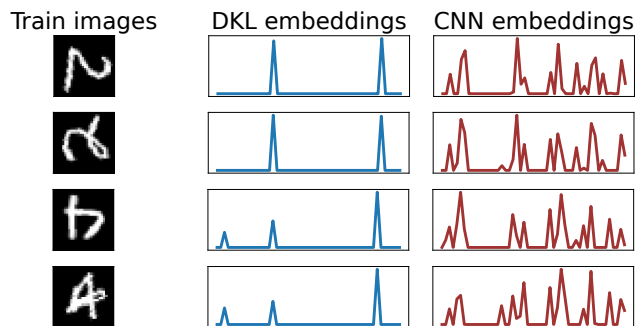


Figure 3: Training images with different orientations and their embeddings. Embeddings produced by joint Deep Kernel Learning (DKL, middle column) are similar for all inputs from one class. Little improvement can be gained on the training data by being rotationally invariant. NN embeddings on the right differ depending on input orientation – signal to learn $p(x_a|x, \phi)$ from.

shortcomings of this approach (Ober et al., 2021; Bradshaw et al., 2017; van Amersfoort et al., 2021): the DKL marginal likelihood correctly penalises complexity for the last layer only, while the NN hyperparameters can still overfit. In our setting, i.e. when trying to combine deep kernel learning with invariance learning, joint training produces overfit weights which results in simplistic features with little intra-class variation⁴. In particular, all training points from the same class are mapped to very similar activations, independent of orientation. This causes a loss of signal for the invariance parameters (see Fig. 3).

use a simple NN with max-likelihood (see Appendix).

⁴This behavior makes sense: The DKL marginal likelihood only penalises complexity in the last layer, (i.e. the GP). The simplistic features from Fig. 3 can be classified by a simple function in the last layer, thus the complexity penalty is small, and the solution has high marg. likelihood.

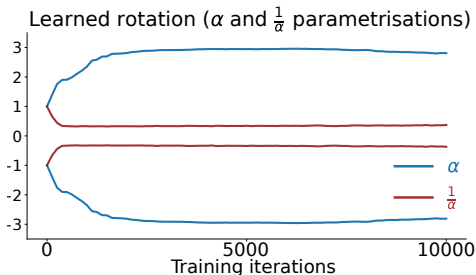


Figure 4: Learned rot. angles parametrised by α and $\frac{1}{\alpha}$. The α -parametrisation, in blue, learns rotational invariance w.r.t. ± 2.8 radians. The $\frac{1}{\alpha}$ -parametrisation (red) learns invariance w.r.t. $\pm \frac{1}{0.37} = \pm 2.7$ radians.

Coordinate ascent training fixes this problem. We pre-train the NN using negative log-likelihood loss. Then, we replace the fully connected last layer with an invariant GP. The marginal likelihood is a good objective given fixed weights (we obtain a GP on transformed inputs), so we fix the NN weights. However, *some* adaptation of the NN to the transformed inputs is beneficial. We thus continue training by alternating between updating the NN, and the GP variational parameters and orbit parameters, hereby successfully learning invariances. (See Fig. 7 and 8: flat parts of the training curves indicate NN training where all kernel hyperparameters, including invariances, remain fixed. When to toggle between the GP and NN training phase is determined using validation data.)

Choosing an invariance parameterisation is simple with our method. Other invariance learning approaches, e.g. Benton et al. (2020) and Schwöbel et al. (2020) rely on explicitly regularising augmentation parameters to be large, and thus require interpretability of their parameters. The marginal likelihood objective is *independent of parameterisation*. To illustrate this we compare parameterising the range of angles by the angle in radians α and by its reciprocal $\xi = \frac{1}{\alpha}$. In the rotMNIST example (see Fig 4) large invariances are needed. This corresponds to large α or small ξ – our method obtains this in both parameterisations. In

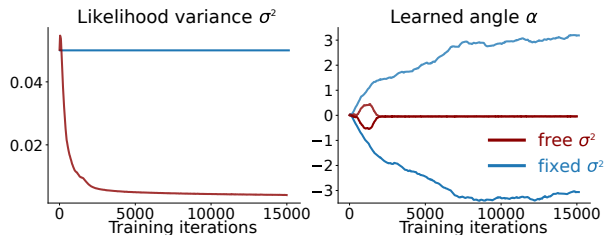


Figure 5: Runs with fixed (red) and non-fixed (blue) kernel and likelihood variance on rotMNIST. The augmentation distribution collapses for non-fixed variances.

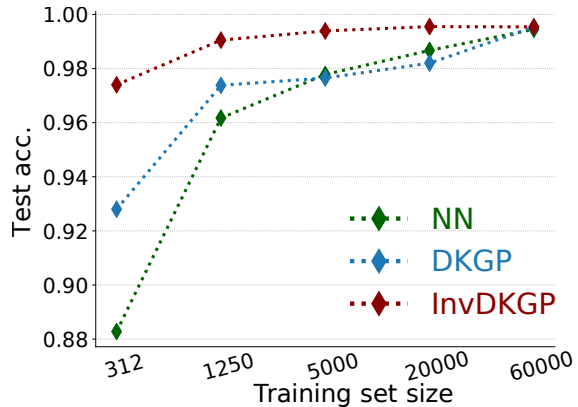


Figure 6: Test accuracies against the training set size on MNIST. We see the invariant model (in red) generalises significantly better, especially for small training sets.

contrast, explicitly regularising invariance parameters to be large would fail for ξ . We wish to stress that generating the orbit distributions is not restricted to affine image transformation and parameterisation independence will be more important as more complicated, non-interpretible invariances are considered.

The Gaussian likelihood is chosen by Van der Wilk et al. (2018) due to its closed-form ELBO. For classification problems, this is a model misspecification. The penalty for not fitting the correct label value becomes large and we can therefore overfit the training data. To alleviate this problem, we fix likelihood and kernel variance (see Fig. 5). The fixed values were determined by trying out a handful candidates – this was sufficient to make invariance learning work. To remove this manual tuning, we will derive an ELBO that works with likelihoods like Softmax in Sec 6.

5.1 MNIST subsets – the low data regime

Having developed a successful training scheme we evaluate it on MNIST subsets. The generalisation problem is particularly difficult when training data is scarce. Inductive biases are especially important and usually parameter-rich neural networks rely on heavy data augmentation when applied to smaller datasets. We train on different subsets of MNIST (LeCun et al.). InvDKGPs outperform both NNs and non-invariant deep kernel GPs. The margin is larger the smaller the training set – with only 1250 training examples we can nearly match the performance of a NN trained on full MNIST (Fig. 6). We conclude it is possible to learn useful invariances even from small data (see Fig. 7). This data efficiency is desirable since models trained on small datasets benefit crucially from augmentation.

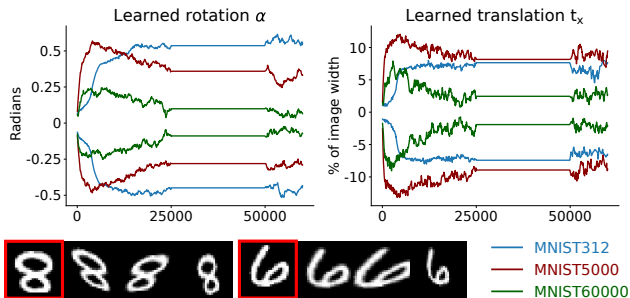


Figure 7: *Top*: Learned invariance parameters (rotation α in radians and x-translation t_x) for a small, medium and large training set. We learn larger α for the smaller subsets. Here, data augmentation is more beneficial. *Bottom*: Two training images x (red frames) and samples from $p(x_a|x, \phi)$ (following columns) learned by the InvDKGP on MNIST using only 312 images.

6 CORRECTING MODEL MISSPECIFICATION

The key observation for inference under the Gaussian likelihood was the *unbiasedness* of the estimators. In this section, we introduce a controlled bias to allow for easy inference in a wide class of likelihoods. In the limit of infinite sampling, the bias disappears and the invariance does not add additional approximation error.

Recall that $f(x)$ constructed in (4) is intractable but can be estimated by Monte Carlo sampling

$$\hat{f}(x) := \frac{1}{S_o} \sum_{i=1}^{S_o} g(x_a^i), \quad (11)$$

where $x_a^i \sim p(x_a|x, \phi)$. Notice,

$$f(x) = \mathbb{E}_{\prod_{i=1}^{S_o} p(x_a^i|x, \phi)} [\hat{f}(x)] =: \tilde{\mathbb{E}} [\hat{f}(x)], \quad (12)$$

where $\prod_{i=1}^{S_o} p(x_a^i|x, \phi)$ is the product density over S_o orbit densities. We remark that f is deterministic in x but stochastic in g , which is a GP. Thus, we can write

$$\mathbb{E}_{q(f(x))} [\log p(y|f(x))] = \mathbb{E}_{q(g)} [\log p(y|f(x))] \quad (13)$$

$$= \mathbb{E}_{q(g)} \left[\log p \left(y \mid \tilde{\mathbb{E}}[\hat{f}(x)] \right) \right] \quad (14)$$

$$\geq \mathbb{E}_{q(g)} \left[\tilde{\mathbb{E}} \left[\log p \left(y \mid \hat{f}(x) \right) \right] \right]. \quad (15)$$

The inequality is due to Jensen’s inequality if the likelihood is *log-concave* in f .⁵ This holds for many common likelihoods, e.g. Gaussian and Softmax.

⁵Nabarro et al. (2021) use this same construction in the weight-space of neural networks to find valid posteriors in the presence of data augmentation, although without invariance learning.

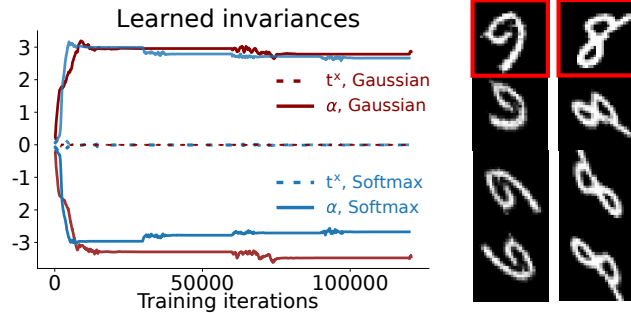


Figure 8: *Left*: Learned invariance parameters (rotation α in radians and x-translation t_x) for rotMNIST. Both the Gaussian and the Softmax model learn to be almost fully rotationally invariant (i.e. $\alpha_{\min/\max} \approx \pm\pi$), and not to be invariant w.r.t. translation (i.e. $t_{\min/\max}^x \approx 0$). Note the different scaling of the y -axis to Fig. 7. *Right*: Two training images (red frames) and samples from orbits.

Equality holds above when $\text{Var}(\hat{f}(x)) = 0$, i.e. the bound becomes tighter as S_o increases (see also Burda et al., 2016). Hence aggressive sampling recovers accurate VI. The right-hand side of (15) can now, without additional bias, be estimated by

$$\frac{1}{S_g} \sum_{k=1}^{S_g} \frac{1}{S_A} \sum_{j=1}^{S_A} \log p \left(y \mid \frac{1}{S_o} \sum_{i=1}^{S_o} g_k(x_a^{ji}) \right). \quad (16)$$

Since extensive sampling is required to keep the bound above tight, it is important to do this efficiently. From a GP perspective this is handled with little effort by sampling the approximate posteriors $q(g)$ using Mathéron’s rule (Wilson et al., 2020). Thus, sampling S_g GPs is cheap compared to sampling from the orbit. S_A denotes the number of \hat{f} samples, this can be fixed to 1 as long as S_o is large.

Summarising, we have shown how we can infer through the marginal likelihood, for the wide class of log-concave likelihoods, by maximising the stochastic ELBO:

$$\mathcal{L} = \frac{1}{S_g} \sum_{k=1}^{S_g} \frac{1}{S_A} \sum_{j=1}^{S_A} \left[\log p \left(y \mid \frac{1}{S_o} \sum_{i=1}^{S_o} g_k(h_w(x_a^{ji})) \right) \right. \\ \left. - \text{KL}[q(\mathbf{u}) \parallel p(\mathbf{u})], \right] \quad (17)$$

$$\text{with } x_a^{ij} \sim p(x_a|x, \phi). \quad (18)$$

The benefits of our new sample based bound are threefold: It broadens model specification, avoids hand-picking and fixing the artificial Gaussian likelihood variance, and doubles training speed.

	Model	Likelihood	Test acc.
M1	NN	Softmax	0.9433
M2	Non-inv. Shallow GP	Gaussian	0.8357
M3	Non-inv. Shallow. GP	Softmax	0.7918
M4	Inv. Shallow GP	Gaussian	0.9516
M5	Inv. Shallow. GP	Softmax	0.9316
M6	Non-inv. Deep Kernel GP	Gaussian	0.9387
M7	Non-inv. Deep Kernel GP	Softmax	0.9351
M8	Inv. Deep Kernel GP	Gaussian	0.9896
M9	Inv. Deep Kernel GP	Softmax	0.9867

Table 1: Test accuracies on rotated MNIST. Invariant models outperform non-invariant counterpart. So do deep kernels contra shallow ones. The invariant deep kernel GPs perform best, outperforming state-of-the-art of 0.989 for learned invariance (Benton et al., 2020).

6.1 Rotated MNIST

The rotated MNIST dataset⁶ was generated from the original MNIST dataset by randomly rotating the images of hand-written digits between 0 and 2π radians. It consists of a training set of 12,000 images along with 50,000 images for testing. We pretrain the neural network from Sec. 5.1 on rotated MNIST (Table 1, M1) and proceed as outlined in Sec. 5. As discussed in Sec. 5, we do not have guarantees that the ELBO acts as a good model selector for the neural network hyperparameters. We thus use a validation set (3000 of the 12000 training points) to find hyperparameters for the NN updates. Once a good training setting is found we re-train on the entire training set (see Appendix for settings). Fig. 8 shows the learned invariances (we use the full ϕ parameterisation but only plot rotation and x -translation for brevity). Both Gaussian and Softmax models learn to be rotation-invariant close the full 2π rotations present in the data. Table 1 contains test accuracies. Deep kernel GPs outperform their shallow counterparts by large margins (differences in test accuracy of ≥ 10 percent points). The same is true for invariant compared to non-invariant models (≥ 3 percent points). While both likelihoods achieve similar test accuracies, we observe a $2.3\times$ speedup per iteration in training for the sample-based Softmax over the Gaussian model. (Gaussian model: 2.64 seconds per iteration, Gaussian + sample bound: 1.32 sec./iter., Softmax + sample bound: 1.13 sec./iter. All runs are executed on 12 GB Nvidia Titan X/Xp GPUs.)

6.2 PatchCamelyon

The PatchCamelyon (PCam, CC0 License, Veeling et al. (2018)) dataset consists of histopathology scans

⁶https://sites.google.com/a/lisa.iro.umontreal.ca/public_static_twiki/variati ons-on-the-mni st-di gi ts

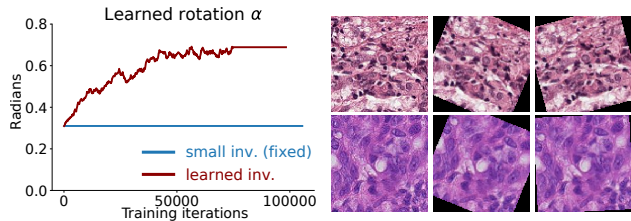


Figure 9: *Left*: Learned rotation on PCam. *Right*: PCam orbit samples. Augmented images look smoother due to interpolation, thus we preprocess the dataset with small rotations when learning invariances.

Model	Test acc.
NN	0.7905
Deep Kernel GP + no inv.	0.8018
NN + small inv.	0.7420
Deep Kernel GP + small inv.	0.8115
Deep Kernel GP + learned inv.	0.8171

Table 2: PCam results. InvDKGP performs best.

of lymph nodes measuring $96 \times 96 \times 3$ pixels. Labels indicate whether the centre patch contains tumor pixels. Veeling et al. (2018) improve test performance from 0.876 to 0.898 by using a NN which is invariant to (hard-coded) 90° rotations of the input. Such discrete, non-differentiable augmentations are not compatible with our backprop-based method, so we instead use continuously sampled rotations (a special case of the transformations described in Sec. 3.3 with $\phi = \alpha$ and $\alpha_{min} = -\alpha_{max}$). This, contrary to Veeling et al. (2018)’s approach, introduces the need for padding and interpolation (see Fig. 9, left), effectively changing the data distribution. We thus apply small rotations as a pre-processing step ($\pi/10$ radians, ‘small inv.’ in Table 2). This lowers performance for a NN alone, i.e. when pre-training. The invariant models counterbalance this performance drop, and the learned invariances produce the best results in our experiments; however, they remain subpar to Veeling et al. (2018). This is due to the limitation to differentiable transformations, as well as our simpler NN (see Appendix). We highlight that our task is fundamentally different: instead of hard-coding invariances we *learn* those during optimisation.

7 EXPLORING LIMITATIONS

Rotated MNIST and PCAM are relatively simple datasets that can be modelled using small NNs. To investigate whether our approach can be used on more complex datasets, we turn to CIFAR-10 (Krizhevsky, 2009), which is usually trained with larger models and data augmentation. Unfortunately, we found that we were unable to learn invariances for CIFAR-10.

To understand why, we designed a simple experiment. We first pretrain ResNet-18-based (He et al., 2016) networks with different levels ν of invariance transformations (see the Appendix for a definition of ν). We then train sparse GP regression (SGPR; Titsias, 2009) models on an augmented training set created by propagating ten points sampled from the augmentation distribution through these neural networks. The samples are generated at different levels of invariance ν , not necessarily matching the levels of the pretrained NNs. We plot the results in Fig. 10: when the network is trained at a small invariance level ϵ , the performance of the SGPR model is highest at an invariance level of 0.01, and rapidly drops off for larger invariances (note the logarithmic x scale). We see a similar result for the network trained at a level of 0.1. Finally, when the network is trained at the same level that the orbit points for the SGPR model are sampled at (‘adapted’), we see that added invariance helps the accuracy, with no steep drop off in accuracy for larger invariances. Therefore, adding invariance does help, but only when the network has already been adapted to that invariance. Currently, in our method this coadaptation is prevented by the current need for coordinate ascent training (Sec. 5).

This experiment indicates that for datasets requiring larger neural networks, we are in a difficult position. We need to adapt the feature extractor jointly with the invariances. However, this approach leads to pathologies as the neural network parameters are not protected from overfitting (Ober et al. (2021), see Sec. 5), which we previously mitigated with coordinate ascent. Therefore, relying on the marginal likelihood to learn invariances with a large feature extractor can easily lead to unwanted behavior – this behavior prevents us from learning these invariances as easily as the marginal likelihood promises. We believe that ongoing research in Bayesian deep learning will alleviate this problem. Bayesian neural networks with methods for marginalising over lower layers too, thus protecting them against overfitting, will render our approach more easily applicable. Such advances will allow us to learn invariances more easily and on more complex tasks than we did for the MNIST and PCAM datasets.

8 CONCLUSION

Neural networks depend on good inductive biases in order to generalise well. Practitioners usually – successfully – handcraft inductive biases, but the idea of learning them from data is appealing. Might we automate the modelling pipeline, moving from hand-crafted models to data driven models; much like we replaced hand-crafted features with learned features in deep neural networks? This work proposes one step

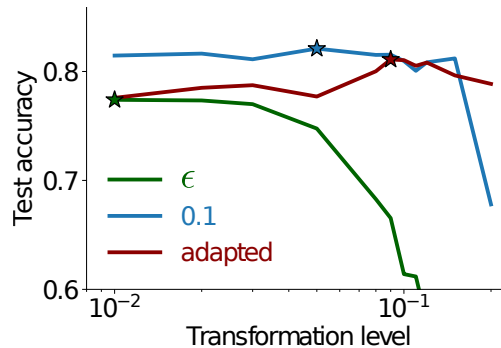


Figure 10: Test accuracies on CIFAR-10 for different transformation levels ν with pretraining at negligible “ ϵ ”, 0.1, and adapted levels. The maxima for each curve are marked with a star, and occur at test accuracies of 77.4% for “ ϵ ”, 82.1% for 0.1, and 81.1% for adapted levels.

in this direction. Inspired by Bayesian model selection we employ the marginal likelihood for learning inductive biases. We avoid the intractability of the marginal likelihood for neural networks by using Deep Kernel Learning. This enables us to leverage previous work on invariance learning in GPs for learning data augmentation in neural networks. We learn useful invariances and improve performance, but encounter challenges when optimising our models. We introduce a new sampling-based bound to the ELBO allowing for inference for the Softmax likelihood, the natural choice for classification tasks, hereby alleviating some of the optimisation difficulties. Others we identify as fundamental limitations of the Bayesian last layer approach.

Societal Impact: This work is situated within basic research in probabilistic ML and, as such, bears all the risks of automation itself: harmful redistribution of wealth to those with access to compute resources and data, loss of jobs, and the environmental impact of such technologies. In fact, our model is more computationally heavy than a standard neural network with hand-tuned data augmentation. However, in the long term, automatic model selection has the potential to *reduce* the need for hyperparameter tuning, which usually dramatically exceeds the resources needed for training the final model.

Acknowledgements

MJ is supported by a research grant from the Carlsberg Foundation (CF20-0370). SWO acknowledges the support of the Gates Cambridge Trust for his doctoral studies.

References

- G. Benton, M. Finzi, P. Izmailov, and A. G. Wilson. Learning invariances in neural networks. In *Advances in Neural Information Processing Systems*, 2020.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 2015.
- J. Bradshaw, A. G. d. G. Matthews, and Z. Ghahramani. Adversarial examples, uncertainty, and transfer testing robustness in Gaussian process hybrid deep networks. *arXiv preprint arXiv:1707.02476*, 2017.
- T. D. Bui, J. Yan, and R. E. Turner. A unifying framework for Gaussian process pseudo-point approximations using power expectation propagation. *Journal of Machine Learning Research (JMLR)*, 18(104):1–72, 2017.
- Y. Burda, R. Grosse, and R. Salakhutdinov. Importance weighted autoencoders. 2016.
- D. R. Burt, C. E. Rasmussen, and M. van der Wilk. Convergence of sparse variational inference in Gaussian processes regression. *Journal of Machine Learning Research (JMLR)*, 21(131):1–63, 2020.
- R. Calandra, J. Peters, C. E. Rasmussen, and M. P. Deisenroth. Manifold Gaussian processes for regression. In *International Joint Conference on Neural Networks (IJCNN)*, 2016.
- E. D. Cubuk, B. Zoph, D. Mane, V. Vasudevan, and Q. V. Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020.
- A. Damianou and N. D. Lawrence. Deep Gaussian processes. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2013.
- T. Dao, A. Gu, A. Ratner, V. Smith, C. De Sa, and C. Ré. A kernel theory of modern data augmentation. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- V. Dutordoir, M. van der Wilk, A. Artemev, and J. Hensman. Bayesian image classification with deep convolutional Gaussian processes. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- V. Dutordoir, J. Hensman, M. van der Wilk, C. H. Ek, Z. Ghahramani, and N. Durrande. Deep neural networks as point estimates for deep Gaussian processes. In *Advances in Neural Information Processing Systems*, 2021.
- D. Ginsbourger, X. Bay, O. Roustant, and L. Carraro. Argumentwise invariant kernels for the approximation of invariant functions. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 21, pages 501–527, 2012.
- D. Ginsbourger, N. Durrande, and O. Roustant. Kernels and designs for modelling invariant functions: From group invariance to additivity. In *mODA 10—Advances in Model-Oriented Design and Analysis*, pages 107–115. Springer, 2013.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- J. Hensman, A. G. d. G. Matthews, and Z. Ghahramani. Scalable variational Gaussian process classification. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- G. E. Hinton and R. Salakhutdinov. Using deep belief nets to learn covariance kernels for Gaussian processes. *Advances in Neural Information Processing Systems*, 2007.
- D. Ho, E. Liang, X. Chen, I. Stoica, and P. Abbeel. Population based augmentation: Efficient learning of augmentation policy schedules. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- A. Immer, M. Bauer, V. Fortuin, G. Rätsch, and M. E. Khan. Scalable marginal likelihood estimation for model selection in deep learning. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR)*, 2015.
- I. R. Kondor. *Group theoretical methods in machine learning*. Columbia University, 2008.
- A. Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- Y. LeCun, C. Cortes, and C. J. Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>.
- J. Lorraine, P. Vicol, and D. Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *Proceedings of the 23rd Conference on*

- International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- D. J. C. MacKay. Model comparison and Occam's razor. *Information Theory, Inference and Learning Algorithms*, pages 343–355, 2003.
- A. G. d. G. Matthews, J. Hensman, R. E. Turner, and Z. Ghahramani. On sparse variational methods and the Kullback-Leibler divergence between stochastic processes. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- A. G. d. G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman. GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research (JMLR)*, 18(40):1–6, 2017.
- M. Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, 1961.
- S. Nabarro, S. Ganev, A. Garriga-Alonso, V. Fortuin, M. van der Wilk, and L. Aitchison. Data augmentation in bayesian neural networks and the cold posterior effect, 2021.
- S. W. Ober and L. Aitchison. Global inducing point variational posteriors for Bayesian neural networks and deep Gaussian processes. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2020.
- S. W. Ober, C. E. Rasmussen, and M. van der Wilk. The promises and pitfalls of deep kernel learning. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2021.
- C. E. Rasmussen and Z. Ghahramani. Occam's razor. *Advances in Neural Information Processing Systems*, 2001.
- P. Schwöbel, F. Warburg, M. Jørgensen, K. H. Madsen, and S. Hauberg. Probabilistic spatial transformers for Bayesian data augmentation. *arXiv preprint arXiv:2004.03637*, 2020.
- M. Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2009.
- R. E. Turner and M. Sahani. Two problems with variational expectation maximisation for time-series models. In D. Barber, T. Cemgil, and S. Chiappa, editors, *Bayesian Time Series Models*, chapter 5, pages 109–130. Cambridge University Press, 2011.
- J. van Amersfoort, L. Smith, A. Jesson, O. Key, and Y. Gal. On feature collapse and deep kernel learning for single forward pass uncertainty. *arXiv preprint arXiv:2102.11409*, 2021.
- M. van der Wilk, M. Bauer, S. T. John, and J. Hensman. Learning invariances using the marginal likelihood. In *Advances in Neural Information Processing Systems*, 2018.
- M. van der Wilk, V. Dutordoir, S. T. John, A. Artemev, V. Adam, and J. Hensman. A framework for inter-domain and multioutput Gaussian processes. *arXiv preprint arXiv:2003.01115*, 2020.
- B. S. Veeling, J. Linmans, J. Winkens, T. Cohen, and M. Welling. Rotation equivariant CNNs for digital pathology. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018.
- C. K. I. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. MIT Press Cambridge, MA, 2006.
- A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Deep kernel learning. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016a.
- A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing. Stochastic variational deep kernel learning. *Advances in Neural Information Processing Systems*, 2016b.
- J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- A. Zhou, T. Knowles, and C. Finn. Meta-learning symmetries by reparameterization. In *9th International Conference on Learning Representations (ICLR)*, 2021.

Supplementary Material: Last Layer Marginal Likelihood for Invariance Learning

A IS THE MARGINAL LIKELIHOOD NECESSARY?

Sec. 1 motivated the marginal likelihood for invariance learning. Given that this loss function is notoriously difficult to evaluate, we verify experimentally that using it is indeed *necessary*, i.e. that the standard maximum likelihood loss is insufficient. Fig. 11 shows invariances learned on rotated MNIST (rotMNIST, see Sec. 6.1 for a description of the dataset) by using a neural network with maximum likelihood loss for two initialisations (blue, green). They collapse as suggested by the theory. The marginal likelihood solution (red) instead identifies appropriate invariances.

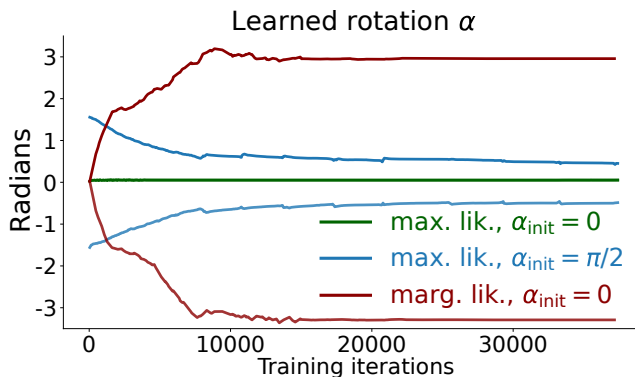


Figure 11: Max. likelihood (green, blue, *collapsing*) and marg. likelihood (red, *useful*) invariances.

B EXPERIMENTAL DETAILS

Exploiting the ideas from Sec. 5, we start by training convolutional neural networks (CNNs, see below for architecture details). After pre-training the CNN, we replace the last fully connected layer with a GP and continue training. In the non-invariant case we train all parameters jointly from here. When learning invariances, we iterate between updating the GP variational- and hyperparameters, and the neural network weights.

B.1 MNIST variations

We here summarise the training setups for the experiments on MNIST variations, i.e. MNIST subsets (Sec. 5.1) and rotated MNIST (Sec. 6.1). We start by outlining the shared neural network architecture and will then list the hyperparameter settings for MNIST and rotMNIST, respectively.

The CNN architecture used in the (rot)MNIST experiments is depicted in Table 3. For rotated MNIST, we train the model for 200 epochs with the Adam optimiser (default parameters). For the MNIST subsets, we train for 60k iterations which corresponds to 200 epochs for the full dataset and respectively more epochs for smaller subsets. The remaining parameters are the same in all experiments: batch size 200, learning rate 0.001, no weight decay, other regularisation or data augmentation. In the pre-training phase we minimise negative log-likelihood, for updates during coordinate ascent we use the ELBO as a loss-function.

Hyperparameter initialisation for the MNIST subset experiments are lengthscale 10, likelihood variance 0.05, kernel variance 1 (fixed likelihood and kernel variance for the invariant model, see Sec. 5.1), posterior variance 0.01. We use 1200 inducing points which we initialise by first passing the images through the neural

Layer	Specifications
Convolution	filters=20, kernel size=(5, 5), padding=same, activation=ReLU
Max pooling	pool size=(2, 2), stride=2
Convolution	filters=50, kernel size=(5, 5), padding=same, activation=ReLU
Max pooling	pool size=(2, 2), stride=2
Fully connected	neurons=500, activation=ReLU
Fully connected	neurons=50, activation=ReLU
Fully connected	neurons=10, activation=Softmax

Table 3: Neural network architecture for MNIST variations. After pre-training, the last fully connected layer (below dashed line) is replaced with a GP layer for the deep kernel models.

network, then using the ‘greedy variance’ method (Burt et al., 2020) on the extracted features. For the smallest dataset MNIST312 we use 312 inducing points only. The batch size is 200 and we choose learning rate 0.001 for the Adam optimiser. For the invariant models, the orbit size is 120 and affine parameters are initialised at $\phi_{min} = \phi_{max} = 0.02$, i.e. we initialise with a small invariance. Without this initialisation we encountered occasional numerical instabilities (Cholesky errors) on the small dataset runs. During coordinate ascent (InvDKGP models) we toggle between training GP and CNN after 25k steps.

Hyperparameter initialisation for the rotMNIST experiment as follows: For all models, we initialise kernel variance 1 (fixed at 1 for $M9$, see Sec. 5.1) and posterior variance 0.01. We use 1200 inducing points. For the invariant models, the orbit size is 120 and affine parameters are initialised at $\phi_{min} = \phi_{max} = 0$, i.e. invariances are learned from scratch. When using coordinate ascent (InvDKGP models, $M8$ & $M9$) we toggle between training GP and CNN after 30k steps. We train different models for a different number of iterations, all until the ELBO has roughly converged. Batch size 200 is used for all models. The remaining initialisations differ between models and are summarised in Table 4.

	Model	Lengthsc.	Lik. var.	LR (decay)
M2	Non-inv. Shallow GP + Gaussian	10	0.02	0.001
M3	Non-inv. Shallow. GP + Softmax	10	-	0.001
M4	Inv. Shallow GP + Gaussian	10	0.05	0.001
M5	Inv. Shallow. GP + Softmax	10	-	0.001
M6	Non-inv. Deep Kernel GP + Gaussian	10	0.05	0.001
M7	Non-inv. Deep Kernel GP + Softmax	20	-	0.001
M8	Inv. Deep Kernel GP + Gaussian	50	0.05 (F)	0.003 (steps / cyclic)
M9	Inv. Deep Kernel GP + Softmax	9	-	0.003 / 0.0003 (s / c)

Table 4: Training settings for rotMNIST models: Kernel lengthscale and likelihood variance initialisations (‘F’ indicates a fixed likelihood variance, see Sec. 5.1). The learning rate column (‘LR’) also indicates whether the learning rate was decayed in the GP/CNN update phases of coordinate ascent. For the ‘steps’(s) decay, we divide by 10 after 50% and again 75% of iterations, for the ‘cyclic’(c) decay, learning rates are: [LR/100, LR/10, LR, LR/10, LR/100]. These training hyperparameters are determined using a validation set (see Sec. 6.1).

B.2 PCam

The CNN architecture is a VGG-like convolutional neural network⁷ described in Table 5. The model is trained for 5 epochs using the Adam optimiser with batch size 64. We use learning rate 0.001 which we divide by 10 after 50% and again 75% of training iterations. In the fully connected block we use dropout with 50% probability when pre-training. Dropout is disabled when training the deep kernel models.

Hyperparameters for the deep kernel GP experiments on PCam are: lengthscale 10 (1 for the learned invariance model), kernel variance 1, posterior variance 0.01. We use 750 inducing points which we initialise as in the previous experiments. The batch size is 32. For PCam we use coordinate ascent for all models since

⁷We closely follow <https://geertliltjens.nl/post/getting-started-with-camel-yon/>.

Layer	Specifications
Convolution	filters=16, kernel size=(3, 3), padding=valid, activation=ReLU
Convolution	filters=16, kernel size=(3, 3), padding=valid, activation=ReLU
Max Pooling	pool size=(2, 2), strides=2
Convolution	filters=32, kernel size=(3, 3), padding=valid, activation=ReLU
Convolution	filters=32, kernel size=(3, 3), padding=valid, activation=ReLU
Max Pooling	pool size=(2, 2), stride=2
Convolution	filters=64, kernel size=(3, 3), padding=valid, activation=ReLU
Convolution	filters=64, kernel size=(3, 3), padding=valid, activation=ReLU
Max Pooling	pool size=(2, 2), stride=2
Fully Connected	neurons=256, activation=ReLU
Dropout	probability=0.5
Fully Connected	neurons=50, activation=None
Dropout	probability=0.5
Fully connected	neurons=2, activation=Softmax

Table 5: Neural network architecture for PCAM. After pre-training, the last fully connected layer (below dashed line) is replaced with a GP layer for the deep kernel models and dropout is disabled.

this improves training stability. Learning rates are 0.001 for the GP update steps and 0.0001 for the CNN updates, no LR decay. We toggle between the two coordinate ascent phases after 50k and 75k iterations in the non-invariant and invariant case, respectively. For the invariant models the orbit size is 20 and we initialise the rotation invariance with $\phi_{min/max} = \alpha_{min/max} = \pm\pi/10$.

B.3 CIFAR-10

Throughout our experiments, we train on a subset of 45,000 points from the full CIFAR-10 (Krizhevsky, 2009) training set and report results on the remaining 5,000 points, as a validation set.

The model we use is a sparse GP regression (SGPR; Titsias (2009)) model with a sum kernel corresponding to a Monte Carlo estimate of the kernel of Eq. 5, using an automatic relevance determination (ARD) squared exponential kernel as a base kernel. We achieve this by sampling 10 points from the full orbit for each data point, and propagating the points through the pretrained feature extractor. For the feature extractor, we choose a ReLU ResNet-18 architecture (He et al., 2016) with an output dimension of 50, using the post-ReLU features. Therefore, for our training set, we end up with a set of $45,000 \times 10 \times 50$ datapoints, where we sum over the 10 orbit samples.

Hyperparameters were chosen as follows. We pretrain the ResNet-18 by adding an additional fully-connected layer with softmax activations. We train for 160 epochs with a batch size of 100 and the Adam optimizer (Kingma and Ba, 2015), starting with a learning rate of 0.001, which we step down by a factor of 10 at epochs 80 and 120. We train the network without weight decay. The SGPR model was subsequently trained for a maximum of 1000 steps, using the Scipy optimizer provided in GPflow (Matthews et al., 2017). During training, we initially set the jitter to $1e-6$, which we increased by a factor of ten if the Cholesky decomposition failed. For the SGPR model, we use 1000 inducing points, initialised as above. We found empirically that the likelihood variance did not have a significant impact on the results; we therefore fixed it to 0.01. Recalling that $\phi = (\alpha, s^x, s^y, p^x, p^y, t^x, t^y)$, we parameterize the transformation by considering the “transformation level” ν such that

$$\phi_{\max} = (0, 1, 1, 0, 0, 0, 0) + \nu \times (\pi, 1, 1, 1, 1, 1, 1), \quad (19)$$

$$\phi_{\min} = (0, 1, 1, 0, 0, 0, 0) - \nu \times (\pi, 1, 1, 1, 1, 1, 1). \quad (20)$$

For the “ ϵ ” setting of the transformation level, we assign $\nu = 0.01$. We chose a non-zero value to ensure that any reduction in performance would be due to a different value of ν , and not because of the lack of presence of the image interpolator in the pretraining (see Sec. 6.2).