
Information Theoretic Meta Learning with Gaussian Processes

Michalis K. Titsias¹

Francisco J. R. Ruiz¹

Sotirios Nikoloutsopoulos²

Alexandre Galashov¹

¹DeepMind

²Athens Univ. of Econ. and Business, Greece

Abstract

We formulate meta learning using information theoretic concepts; namely, mutual information and the information bottleneck. The idea is to learn a stochastic representation or encoding of the task description, given by a training set, that is highly informative about predicting the validation set. By making use of variational approximations to the mutual information, we derive a general and tractable framework for meta learning. This framework unifies existing gradient-based algorithms and also allows us to derive new algorithms. In particular, we develop a memory-based algorithm that uses Gaussian processes to obtain non-parametric encoding representations. We demonstrate our method on a few-shot regression problem and on four few-shot classification problems, obtaining competitive accuracy when compared to existing baselines.

1 INTRODUCTION

Meta learning [Ravi and Larochelle, 2017, Vinyals et al., 2016, Edwards and Storkey, 2017, Finn et al., 2017, Lacoste et al., 2019, Nichol et al., 2018] and few-shot learning [Li et al., 2006, Lake et al., 2011] aim at deriving data-efficient learning algorithms that can rapidly adapt to new tasks. To achieve that, these algorithms train deep neural networks on a set of tasks drawn from a common distribution. Each task is typically divided into a training (or support) set and a validation (or target) set. The neural network is fitted to predict each task’s validation set from its support set. By sharing information across tasks, the network learns to rapidly adapt to new tasks and generalize from few examples at test time.

There is a plethora of work on few-shot learning algorithms, including memory-based [Vinyals et al., 2016, Ravi and Larochelle, 2017] and gradient-based [Finn et al., 2017, Nichol et al., 2018] procedures. Among those, the gradient-

based model agnostic meta learning (MAML) by Finn et al. [2017] has been particularly influential. However, despite the success of these algorithms, meta learning still lacks unifying principles that allow us to relate all these approaches and invent new schemes. While there exist probabilistic interpretations of existing methods, such as the approximate Bayesian inference approach [Grant et al., 2018, Finn et al., 2018, Yoon et al., 2018] and the related conditional probability modelling approach [Garnelo et al., 2018, Gordon et al., 2019], there is not a general and tractable learning principle for meta learning that can help to get a better understanding of existing algorithms and derive new ones.

We address this issue in this paper. Specifically, we introduce an information theoretic view of meta learning, and we derive a general (and practical) framework for meta learning by exploiting variational approximations of the information theoretic quantities. We show that this framework recovers gradient-based algorithms, such as MAML and its probabilistic interpretations, and that it also allows us to derive new methods: we use it to introduce a new memory-based algorithm for supervised few-shot learning.

More in detail, we consider the information bottleneck principle [Tishby et al., 1999], which can learn a stochastic encoding of the support set of each task that is highly informative for predicting the validation set of that task. The stochastic encoding is optimized through the difference between two mutual informations [Cover and Thomas, 2006], so that the encoding compresses the training set into a representation that predicts well the validation set. Since the mutual information is intractable, we exploit recent approximations of the information bottleneck [Alemi et al., 2017, Chalk et al., 2016, Achille and Soatto, 2016] that make use of variational bounds [Barber and Agakov, 2003]; this results in a tractable objective function for meta learning called variational information bottleneck (VIB).

We show that VIB gives rise to gradient-based meta learning methods, such as MAML, when the encoding is *parametric*, i.e., it has some model parameters or neural network weights.

Additionally, VIB recovers a probabilistic version of MAML simply by using a stochastic encoding. Furthermore, we use VIB to develop a memory-based algorithm for supervised few-shot learning (right panel in Figure 1), based on Gaussian processes (GPs) [Rasmussen and Williams, 2006] and deep neural kernels [Wilson et al., 2016], that offers a kernel-based Bayesian view of a memory system. With GPs, the underlying encoding takes the form of a *non-parametric* function that follows a stochastic process amortized by the training set. Our framework is general and would naturally allow for other extensions, such as combinations of memory and gradient-based meta learning.

We demonstrate our GP-based method on a few-shot regression problem and four classification problems. The former is a few-shot sinusoid regression problem, where we obtain smaller error than MAML. For the latter, we use the unified training and evaluation protocol of Patacchiola et al. [2020], obtaining that the GP-based algorithm provides competitive accuracy when compared to existing baselines, obtaining state-of-the-art results in some cases.

2 INFORMATION BOTTLENECK VIEW OF META LEARNING

We wish to learn from a distribution of tasks. During training, we observe a set of tasks, each consisting of a *task description* represented by the support or training set D^t and a *task validation* represented by the target or validation set D^v . At test time, we only have access to the training set D^t of a new task, and the goal of the algorithm is to adapt and form predictions on that task’s D^v or on further test data.

We formulate meta learning using information theoretic concepts. The idea is to learn a stochastic representation or encoding of the task description D^t that is highly informative for predicting D^v . We introduce a random variable Z that represents the encoding; it is drawn from a distribution $q_w(Z|D^t)$ parameterized by w . Thus, the joint distribution over D^t , D^v , and Z is

$$q_w(D^v, D^t, Z) = q_w(Z|D^t)p(D^v, D^t), \quad (1)$$

where $p(D^v, D^t)$ denotes the unknown data distribution.

The goal of meta learning under this view is to tune the parameters w of the encoder. To that end, one approach is to maximize the mutual information $I(Z, D^v)$ between Z and the target set D^v . A trivial way to obtain a maximally informative representation is to set $Z = D^t$, for which $I(Z, D^v)$ attains its maximum value $I(D^t, D^v)$. However, this is not a useful representation since Z simply memorizes the training set and there is no learning involved. Instead, the information bottleneck (IB) principle [Tishby et al., 1999] adds a model complexity penalty to the maximization of $I(Z, D^v)$ to promote a more compressive or parsimonious encoding of the training set D^t , i.e., an encoding that extracts only relevant

information from D^t for predicting D^v . The penalty term is the mutual information $I(Z, D^t)$ that we want to minimize. Thus, the IB objective is

$$\mathcal{L}_{\text{IB}}(w) = I(Z, D^v) - \beta I(Z, D^t), \quad (2)$$

where $\beta \geq 0$ is a hyperparameter. In this IB objective, the mutual information $I(Z, D^t)$ acts as a regularizer in the maximization of $I(Z, D^v)$, and it tries to introduce some partial independence between Z and D^t so that irrelevant information is removed when encoding D^t into Z .

However, the IB objective is intractable because both mutual information terms depend on the unknown data distribution $p(D^v, D^t)$. We overcome this by using variational bounds of the mutual information; in particular we obtain a tractable lower bound on $\mathcal{L}_{\text{IB}}(w)$ by lower bounding $I(Z, D^v)$ and upper bounding $I(Z, D^t)$. The approach is similar to the one by Alemi et al. [2017], who introduced VIB for supervised learning of a single task.

2.1 VARIATIONAL INFORMATION BOTTLENECK (VIB) FOR META LEARNING

Here we construct a lower bound $\mathcal{F} \leq \mathcal{L}_{\text{IB}}(w)$. We first lower bound the mutual information $I(Z, D^v) = \text{KL}[q_w(Z, D^v)||q_w(Z)p(D^v)]$, given by

$$I(Z, D^v) = \mathbb{E}_{q_w(Z, D^v)} \left[\log \frac{q_w(D^v|Z)}{p(D^v)} \right], \quad (3)$$

where KL denotes the Kullback-Leibler divergence and $q_w(D^v|Z) = \frac{\int q_w(Z|D^t)p(D^v, D^t)dD^t}{\int q_w(Z|D^t)p(D^v, D^t)dD^t dD^v}$ is intractable because it involves the data distribution $p(D^v, D^t)$. To lower bound $I(Z, D^v)$, we follow Barber and Agakov [2003] (see Appendix A.1) and introduce a *decoder* model $p_\theta(D^v|Z)$ to approximate the intractable $q_w(D^v|Z)$, where θ are additional parameters, yielding the bound¹

$$\begin{aligned} I(Z, D^v) &\geq \mathbb{E}_{q_w(Z, D^v)} \left[\log \frac{p_\theta(D^v|Z)}{p(D^v)} \right] \\ &= \mathbb{E}_{q_w(Z, D^v)} [\log p_\theta(D^v|Z)] + \mathcal{H}(D^v), \end{aligned} \quad (4)$$

where the entropy $\mathcal{H}(D^v)$ is a constant that does not depend on the tunable parameters (θ, w) . Secondly, we upper bound the mutual information $I(Z, D^t)$,

$$I(Z, D^t) = \mathbb{E}_{q_w(Z, D^t)} \left[\log \frac{q_w(Z|D^t)}{q_w(Z)} \right], \quad (5)$$

where $q_w(Z) = \int q_w(Z|D^t)p(D^t)dD^t$ is intractable because it involves the unknown data distribution $p(D^t)$. Similarly, we approximate $q_w(Z)$ with a tractable *prior* model

¹The bounds are valid even when the parameters w of the encoder $q_w(Z|D^t)$ and θ of the decoder $p_\theta(D^v|Z)$ and prior $p_\theta(Z)$ have shared components (e.g., shared neural network weights).




Figure 1: **(Left)** Meta learning with the information bottleneck. The encoder $q_w(Z_i|D_i^t)$ learns to compress each task’s training set D_i^t by minimizing the mutual information $I(Z_i, D_i^t)$ and to predict the task’s validation set D_i^v by maximizing $I(Z_i, D_i^v)$ (Section 2). **(Right)** Particularization to supervised few-shot learning, where we have input-output data for each task. MAML and GP memory-based methods (developed in this paper) are instances of this framework (Sections 2.2 and 3).

distribution $p_\theta(Z)$, leading to the bound

$$I(Z, D^t) \leq \mathbb{E}_{q_w(Z, D^t)} \left[\log \frac{q_w(Z|D^t)}{p_\theta(Z)} \right]. \quad (6)$$

We obtain the overall bound, $\mathcal{F}(\theta, w) + \mathcal{H}(D^v) \leq \mathcal{L}_{\text{IB}}(w)$, by applying Eqs. 4 and 6 into Eq. 2,

$$\begin{aligned} \mathcal{F}(\theta, w) &= \mathbb{E}_{q_w(Z, D^v)} [\log p_\theta(D^v|Z)] \\ &\quad - \beta \mathbb{E}_{q_w(Z, D^t)} \left[\log \frac{q_w(Z|D^t)}{p_\theta(Z)} \right], \end{aligned} \quad (7)$$

where we have dropped the constant $\mathcal{H}(D^v)$ from the objective function $\mathcal{F}(\theta, w)$.

Thus, given a set of b tasks $\{D_i^t, D_i^v\}_{i=1}^b$, learning the parameters (θ, w) during meta training reduces to maximizing the empirical average $\frac{1}{b} \sum_i \tilde{\mathcal{F}}_i(\theta, w)$, where each $\tilde{\mathcal{F}}_i(\theta, w)$ is an unbiased estimate of $\mathcal{F}(\theta, w)$ based on the i -th task, given by (see Appendix A.2)

$$\begin{aligned} \tilde{\mathcal{F}}_i(w, \theta) &= \mathbb{E}_{q_w(Z_i|D_i^t)} [\log p_\theta(D_i^v|Z_i)] \\ &\quad - \beta \text{KL}[q_w(Z_i|D_i^t) || p_\theta(Z_i)]. \end{aligned} \quad (8)$$

The meta-training procedure is carried out in different episodes, where at each step we receive a minibatch of tasks and perform a stochastic gradient maximization step based on the data in that minibatch.

Note that the objective in Eq. 8 is similar to the variational inference objectives for meta learning [Ravi and Beatson, 2019]. In particular, it can be viewed as an evidence lower bound (ELBO) on the log marginal likelihood of the validation set, $\log \int p_\theta(D_i^v|Z_i) p_\theta(Z_i) dZ_i$, with the following differences: (i) there is a hyperparameter β in front of the KL term, and (ii) the distribution $q_w(Z_i|D_i^t)$ is more restricted than in variational inference, since it now acts as a *stochastic bottleneck* that encodes the support set D_i^t (i.e., it is amortized by D_i^t) and via the term $\mathbb{E}_{q_w(Z_i|D_i^t)} [\log p_\theta(D_i^v|Z_i)]$ it is optimized to reconstruct the validation set.

2.2 INFORMATION THEORETIC VIEW OF MAML-TYPE METHODS

We recover MAML [Finn et al., 2017] as a special case of the VIB framework. To see this, suppose that the encoding variable Z_i for the i -th task coincides with a vector of some task-specific model parameters or neural network weights ψ_i , so that $p_\theta(D_i^v|Z_i) \equiv p(D_i^v|\psi_i)$, and $p_\theta(Z_i) \equiv p_\theta(\psi_i)$ is the prior over these parameters. MAML tries to find a shared initial parameter value θ so that few gradient steps based on the support set objective, $\log p(D_i^t|\theta)$, lead to a task-specific parameter value ψ_i with good predictive capacity on the validation set. That is, MAML estimates the task parameters by $\psi_i = \theta + \Delta(\theta, D_i^t)$, where $\Delta(\theta, D_i^t)$ denotes these inner loop adaptation steps—for one step of stochastic gradient descent (SGD) with step size ρ , it is just $\rho \nabla_\theta \log p(D_i^t|\theta)$. We recover MAML from Eq. 8 by setting $\beta = 0$ and setting the encoder to a deterministic Dirac delta measure, $q_\theta(\psi_i|D_i^t) = \delta_{\theta + \Delta(\theta, D_i^t)}(\psi_i)$, so that $\tilde{\mathcal{F}}_i(\theta) = \log p(D_i^v|\theta + \Delta(\theta, D_i^t))$.

Bayesian or probabilistic MAML. Based on this view, we can generalize MAML by using a probabilistic encoder instead of a Dirac delta. For instance, we can use a Gaussian encoder, $q_{\theta, s}(\psi_i|D_i^t) = \mathcal{N}(\psi_i|\theta + \Delta(\theta, D_i^t), s)$, where s is a diagonal covariance. Then, the objective becomes

$$\begin{aligned} \tilde{\mathcal{F}}_i(\theta, s) &= \mathbb{E}_{\mathcal{N}(\epsilon|0, I)} [\log p(D^v|\theta + \Delta(\theta, D_i^t) + \sqrt{s} \circ \epsilon)] \\ &\quad - \beta \text{KL}[q_{\theta, s}(\psi_i|D_i^t) || p_\theta(\psi_i)], \end{aligned} \quad (9)$$

where we have reparameterized the expectation following Kingma and Welling [2014] for stochastic optimization of the meta parameters (θ, s) . This connects with several approaches in the literature that have introduced probabilistic or Bayesian MAML algorithms [Grant et al., 2018, Finn et al., 2018, Yoon et al., 2018, Nguyen et al., 2019]. From the VIB perspective, such probabilistic MAML methods are associated with a probabilistic encoder that introduces uncertainty, where the hyperparameter β in the VIB objective controls the amount of uncertainty.

2.3 VIB FOR SUPERVISED META LEARNING

Here, we explain how to adapt the VIB principle to supervised meta learning. Consider a few-shot supervised learning problem, where for each task we wish to predict outputs or labels given the corresponding inputs. In this supervised setting, we denote the task support set as $D^t = (Y^t, X^t)$, where $Y^t = \{y_j^t\}_{j=1}^{n^t}$ and $X^t = \{x_j^t\}_{j=1}^{n^t}$ denote the output and input observations, respectively. Similarly, we write $D^v = (Y^v, X^v)$ for the validation set. During meta testing, for any new task we observe the support set $D_*^t = (Y_*^t, X_*^t)$ together with the test inputs X_*^v and the goal is to predict the test outputs Y_*^v .

This suggests that we can construct a task encoder distribution of the form $q_w(Z|Y^t, X)$ that depends on the training outputs Y^t and generally on all inputs $X = (X^t, X^v)$.² We would like to train this encoder so that Z becomes highly predictive about the validation outputs Y^v and simultaneously compressive about Y^t . Then, we form a VIB objective based on the *input-conditioned* information bottleneck, $I(Z, Y^v|X) - \beta I(Z, Y^t|X)$, i.e., where both mutual information terms are conditional. Similarly to Section 2.1, we lower bound this objective and approximate it by an unbiased empirical average, $\frac{1}{b} \sum_{i=1}^b \tilde{\mathcal{F}}_i(\theta, w)$, where

$$\begin{aligned} \tilde{\mathcal{F}}_i(\theta, w) = & \mathbb{E}_{q_w(Z_i|Y_i^t, X_i)} [\log p_\theta(Y_i^v|Z_i, X_i)] \\ & - \beta \text{KL} [q_w(Z_i|Y_i^t, X_i) || p_\theta(Z_i|X_i)]. \end{aligned} \quad (10)$$

Here, $p_\theta(Y_i^v|X_i^v, Z_i)$ and $p_\theta(Z_i|X_i)$ are the decoder and prior model distributions introduced by the variational approximation (see Appendix A.3 for a detailed derivation).

Eq. 10 provides the VIB objective for supervised meta learning. We can recover the supervised version of MAML as a special case of VIB, similarly to Section 2.2.

3 SUPERVISED META LEARNING WITH GAUSSIAN PROCESSES

In this section, we use the VIB framework for meta learning to develop a new algorithm for few-shot supervised learning, which uses a non-parametric stochastic encoder based on a GP model. In Section 3.1 we outline the structure of the GP meta learning method and in Section 3.2 we give further details about how to set up the GP encoder.

3.1 GAUSSIAN PROCESS VIB-BASED METHOD

As described in Section 2.3, the VIB framework for few-shot supervised learning requires us to specify the encoding variable Z_i , together with the encoder $q_w(Z_i|Y_i^t, X_i)$, the

²Dependence on all inputs allows explaining both transductive and non-transductive settings [Bronskill et al., 2020, Finn et al., 2017, Nichol et al., 2018] as special cases (see Appendix B).

decoder over the validation outputs $p_\theta(Y_i^v|Z_i, X_i)$, and the prior model $p_\theta(Z_i|X_i)$. Here, we construct these quantities using a GP model [Rasmussen and Williams, 2006] in order to obtain flexible non-parametric stochastic functions.

GP specification. We introduce an (unknown) task-specific function $f_i(x)$ that is *a priori* (before observing any task data) drawn from a GP, i.e., $f_i(x) \sim GP(0, k_\theta(x, x'))$, where k_θ denotes the kernel function. Without loss of generality, we use a deep kernel function, $k_\theta(x, x') = \sigma_f^2 \phi(x; \theta)^\top \phi(x'; \theta)$, where $\phi(x; \theta)$ is a feature vector given by a deep neural network parameterized by θ , and σ_f^2 is the kernel variance parameter (which, if learnable, we also consider to be part of the full set of parameters θ). Equivalently, we can interpret this construction as setting $f_i(x) = \phi(x; \theta)^\top \theta_i^{\text{out}}$, i.e., a linear function of the feature vector with task-specific weights drawn from a Gaussian distribution, $\theta_i^{\text{out}} \sim \mathcal{N}(\theta_i^{\text{out}}|0, \sigma_f^2 I)$. In the GP formulation, the weights θ_i^{out} are marginalized out and we are left with the parameters θ shared across tasks.

Suppose now that we observe the task data, i.e., the support $D_i^t = (Y_i^t, X_i^t)$ and validation $D_i^v = (Y_i^v, X_i^v)$ sets, so that we can evaluate the task function $f_i(\cdot)$ on all task inputs $X_i = (X_i^t, X_i^v)$. Let $f_{i,j}^v \equiv f(x_{i,j}^v)$ denote the function value at the validation input $x_{i,j}^v$ associated with output $y_{i,j}^v$, where the index j runs over the instances in D_i^v . Let $\mathbf{f}_i^v = \{f_{i,j}^v\}_{j=1}^{n^v}$ be the vector containing all such values. Similarly, let \mathbf{f}_i^t be the vector of function values at the inputs X_i^t .

Components of the VIB formulation. We now specify the four ingredients of VIB: (i) the encoding variable Z_i , (ii) the prior $p_\theta(Z_i|X_i)$, (iii) the decoder $p_\theta(Y_i^v|Z_i, X_i)$, and (iv) the encoder $q_w(Z_i|Y_i^t, X_i)$. (i) We set the task encoding variable Z_i to the full set of function values, $Z_i \equiv (\mathbf{f}_i^v, \mathbf{f}_i^t)$. It is a *non-parametric* encoding, since its size grows with the number of task data points. (ii) We set the prior model to the GP prior, $p_\theta(Z_i|X_i) \equiv p(\mathbf{f}_i^v, \mathbf{f}_i^t|X_i)$, where³

$$\begin{aligned} p(\mathbf{f}_i^v, \mathbf{f}_i^t|X_i) = & p(\mathbf{f}_i^v|\mathbf{f}_i^t, X_i) \times p(\mathbf{f}_i^t|X_i^t), \\ = & \mathcal{N}(\mathbf{f}_i^v|\mathbf{K}_i^{vt}[\mathbf{K}_i^t]^{-1}\mathbf{f}_i^t, \mathbf{K}_i^v - \mathbf{K}_i^{vt}[\mathbf{K}_i^t]^{-1}[\mathbf{K}_i^{vt}]^\top) \\ & \times \mathcal{N}(\mathbf{f}_i^t|\mathbf{0}, \mathbf{K}_i^t). \end{aligned} \quad (11)$$

Here, \mathbf{K}_i^t and \mathbf{K}_i^v are $n^t \times n^t$ and $n^v \times n^v$ kernel matrices on the training X_i^t and validation inputs X_i^v , respectively, and \mathbf{K}_i^{vt} is the $n^v \times n^t$ cross kernel matrix between the two sets of inputs. (iii) We set the decoder model $p_\theta(Y_i^v|\mathbf{f}_i^v, \mathbf{f}_i^t, X_i)$ to the standard GP likelihood. For i.i.d. observations, Y_i^v is independent of \mathbf{f}_i^t and X_i given \mathbf{f}_i^v , and the likelihood factorizes across data points, $p(Y_i^v|\mathbf{f}_i^v) = \prod_{j=1}^{n^v} p(y_{i,j}^v|f_{i,j}^v)$. Each $p(y_{i,j}^v|f_{i,j}^v)$ is a standard likelihood model, such as a Gaussian density $p(y_{i,j}^v|f_{i,j}^v) = \mathcal{N}(y_{i,j}^v|f_{i,j}^v, \sigma^2)$ for regression problems or a categorical/softmax likelihood for few-shot classification (see Appendix C.3). (iv) We set the encoder

³We drop the dependency on θ for notational simplicity.

$q_w(Z_i|Y_i^t, X_i) \equiv q(\mathbf{f}_i^v, \mathbf{f}_i^t|Y_i^t, X_i)$ as the product

$$q(\mathbf{f}_i^v, \mathbf{f}_i^t|Y_i^t, X_i) = p(\mathbf{f}_i^v|\mathbf{f}_i^t, X_i)q(\mathbf{f}_i^t|D_i^t), \quad (12)$$

where $p(\mathbf{f}_i^v|\mathbf{f}_i^t|X_i)$ is the conditional GP prior from Eq. 11, and $q(\mathbf{f}_i^t|D_i^t)$ is an encoder of the training set that takes the form of a Gaussian distribution amortized by D_i^t ; see Section 3.2 for details. Eq. 12 shares a similar structure with a standard posterior GP, where we first observe the training set D_i^t , then we compute the (approximate) posterior $q(\mathbf{f}_i^t|D_i^t)$, and finally we predict the validation set function values at inputs X_i^v based on the conditional GP prior $p(\mathbf{f}_i^v|\mathbf{f}_i^t, X_i)$.

Objective function. Putting all together, we obtain the following VIB single-task objective (see Appendix C.1),

$$\sum_{j=1}^{n^v} \mathbb{E}_{q(f_{i,j}^v)} [\log p(y_{i,j}^v|f_{i,j}^v)] - \beta \text{KL}[q(\mathbf{f}_i^t|D_i^t)||p(\mathbf{f}_i^t|X_i^t)], \quad (13)$$

where $q(f_{i,j}^v) = \int p(f_{i,j}^v|\mathbf{f}_i^t, x_{i,j}^v, X_i^t)q(\mathbf{f}_i^t|D_i^t)d\mathbf{f}_i^t$ is a univariate Gaussian distribution over one validation function value $f_{i,j}^v$. The distribution $q(f_{i,j}^v)$ depends on the training set and the validation input $x_{i,j}^v$. Thus, to compute it, we simply use the univariate conditional GP $p(f_{i,j}^v|\mathbf{f}_i^t, x_{i,j}^v, X_i^t)$ to extrapolate the function value at $x_{i,j}^v$.

3.2 GP ENCODER

We now specify the encoder $q(\mathbf{f}^t|D^t)$ used in Eq. 12. A suitable choice is to set it equal to the exact posterior distribution over \mathbf{f}^t given the training set, i.e., $q(\mathbf{f}^t|D^t) = p(\mathbf{f}^t|D^t) \propto \prod_{j=1}^{n^t} p(y_j^t|f_j^t)\mathcal{N}(\mathbf{f}^t|0, \mathbf{K}^t)$.

The posterior $p(\mathbf{f}^t|D^t)$ is tractable for standard regression problems with Gaussian likelihood, i.e., $p(y_j^t|f_j^t) = \mathcal{N}(y_j^t|f_j^t, \sigma^2)$. In this case, the posterior is also Gaussian,

$$p(\mathbf{f}^t|D^t) = \mathcal{N}(\mathbf{f}^t|\mathbf{K}^t(\mathbf{K}^t + \sigma^2 I)^{-1}Y^t, \mathbf{K}^t - \mathbf{K}^t(\mathbf{K}^t + \sigma^2 I)^{-1}\mathbf{K}^t), \quad (14)$$

and thus the encoder $q(\mathbf{f}^t|D^t) = p(\mathbf{f}^t|D^t)$ is tractable. Note that $q(\mathbf{f}^t|D^t)$ depends on D^t through the kernel matrix \mathbf{K}^t , which mixes all the training inputs in X^t . The posterior covariance depends only on X^t , while its mean depends additionally on Y^t . This choice of the encoder does not require to introduce any extra variational parameters w , as $q(\mathbf{f}^t|D^t)$ depends only on the model parameters θ that appear in the kernel function and (possibly) in the likelihood.

When the likelihood is not Gaussian, the posterior $p(\mathbf{f}^t|D^t)$ is not available analytically. In this case, we set $q(\mathbf{f}^t|D^t)$ to an approximate posterior. Specifically, we approximate each non-Gaussian likelihood term $p(y_j^t|f_j^t)$ with a Gaussian term; this is similar to the Gaussian approximations of variational Bayes or expectation-propagation in GPs [Rasmussen and Williams, 2006, Oppen and Archambeau, 2009,

Hensman et al., 2014]. That is, we approximate

$$p(y_j^t|f_j^t) \approx \mathcal{N}(m_j^t|f_j^t, s_j^t), \quad (15)$$

where $m_j^t \equiv m_w(y_j^t, x_j^t) \in \mathbb{R}$ and $s_j^t \equiv s_w(y_j^t, x_j^t) \in \mathbb{R}_+$ are amortized functions parameterized by w that take as input a data point (y_j^t, x_j^t) associated with the latent variable f_j^t and output the parameters of the Gaussian approximation. In this case, the amortized encoder becomes a fully dependent multivariate Gaussian distribution of the form

$$q(\mathbf{f}^t|D^t) = \mathcal{N}(\mathbf{f}^t|\mathbf{K}^t(\mathbf{K}^t + \mathbf{S}^t)^{-1}\mathbf{m}^t, \mathbf{K}^t - \mathbf{K}^t(\mathbf{K}^t + \mathbf{S}^t)^{-1}\mathbf{K}^t),$$

where \mathbf{S}^t is a diagonal covariance matrix with the vector $(s_1^t, \dots, s_{n^t}^t)$ in its diagonal, and \mathbf{m}^t is the vector of values $(m_1^t, \dots, m_{n^t}^t)$. Using this encoder, we can re-write the VIB objective from Eq. 13 in a form that is computationally more convenient (see Appendix C.1),

$$\sum_{j=1}^{n^v} \mathbb{E}_{q(f_j^v)} [\log p(y_j^v|f_j^v)] - \beta \sum_{j=1}^{n^t} \mathbb{E}_{q(f_j^t)} [\log \mathcal{N}(m_j^t|f_j^t, s_j^t)] + \beta \log \mathcal{N}(\mathbf{m}^t|0, \mathbf{K}^t + \mathbf{S}^t), \quad (16)$$

where each marginal Gaussian distribution $q(f_j)$ is computed using the same expression, $q(f_j) = \mathcal{N}(f_j|\mathbf{k}_j^t(\mathbf{K}^t + \mathbf{S}^t)^{-1}\mathbf{m}^t, k_j - \mathbf{k}_j^t(\mathbf{K}^t + \mathbf{S}^t)^{-1}\mathbf{k}_j^t)$, regardless of whether x_j is from the validation or the training set (or any other further test set). Here, $\mathbf{k}_j^t \equiv k(x_j, X^t)$ is the n^t -dimensional row vector of kernel values between x_j and the training inputs X^t , and $k_j \equiv k(x_j, x_j)$.

We can simplify the encoder by assuming a constant value $s_w(x_j^t) := \sigma^2$ and further assuming that $m_w(y_j^t, x_j^t)$ depends only on the output y_j^t , i.e., $m_w(y_j^t, x_j^t) := m_w(y_j^t)$. In our experiments, we found that this simplification worked better in most few-shot settings.

GP encoder for classification problems. We now particularize the encoder for classification problems, which is the standard application in few-shot learning. For notational simplicity, we focus on binary classification, and we describe multi-class classification in Appendix C.3.

Consider a meta learning problem in which each task is a binary classification problem. The class labels take value in $\{-1, 1\}$ and the (non-Gaussian) likelihood is the sigmoid $p(y_j^t|f_j^t) = 1/(1 + e^{-y_j^t f_j^t})$. To specify the GP encoder, we simply need to choose the form of the amortized functions $(m_w(y_j^t, x_j^t), s_w(y_j^t, x_j^t))$. We set them as

$$m(y_j^t, x_j^t) := y_j^t \times \tilde{m}_w(x_j^t), \quad s_w(y_j^t, x_j^t) := s_w(x_j^t),$$

where $\tilde{m}_w(x_j^t)$ and $s_w(x_j^t)$ are parametrized by a neural network. Note that the dependence on the output label $y_j^t \in \{-1, 1\}$ simply changes the sign of $\tilde{m}_w(x_j^t)$, which yields amortization invariance to class re-labeling. That is, if we

swap the roles of the two labels $\{-1, 1\}$, the amortization remains valid and it does not require any change.

We now describe the simplified encoder. We obtain it by assuming a constant form for the functions $(\tilde{m}_w(x_j^t), s_w(x_j^t)) := (\tilde{m}, \sigma^2)$; this has significant computational benefits because it only requires to tune two scalar parameters (\tilde{m}, σ^2) . As mentioned above, the simplified encoder led to improved performance in our experiments.

4 RELATED WORK

In this work, we formulated meta learning using the VIB principle. VIB has been used before for different purposes, such as regularization of single-task supervised learning [Alemi et al., 2017], sparse coding [Chalk et al., 2016], re-interpretation of variational autoencoders [Burgess et al., 2018] and dropout [Achille and Soatto, 2016], and for compression of deep neural networks [Dai et al., 2018].

A meta learning method that connects with the information bottleneck was recently proposed by Hu et al. [2020]. They use the information bottleneck to analyze the generalization of a variational Bayesian inference objective suitable for transductive supervised few-shot learning. The information bottleneck from Theorem 1 of Hu et al. [2020] differs from the information bottleneck objective in our paper (the objective in Eq. 10 for the supervised learning case) in the second term. This second term is the conditional entropy $\mathcal{H}(Z|D^t) = -E_{q(Z, D^t)}[\log q(Z|D^t)]$ in Eq. 19 of Hu et al. [2020], while it is a mutual information in our case, $I(Z, D^t) = \mathcal{H}(Z) - \mathcal{H}(Z|D^t)$. When used as a regularizer, the conditional entropy $\mathcal{H}(Z|D^t)$ alone may lead to simply increasing the variance of $q(Z|D^t)$ and not extracting the relevant information from D^t to predict D^v . Moreover, this term is a differential entropy and thus it can be unbounded. In contrast, $I(Z, D^t) \geq 0$ is bounded, so it does not allow the variance of $q(Z|D^t)$ to blow and it encourages extracting relevant information from D^t to predict D^v .

Given the probabilistic nature of our framework, we can relate it to other probabilistic or Bayesian approaches, and particularly with those that: (i) probabilistically re-interpret and extend gradient-based methods [Grant et al., 2018, Finn et al., 2018, Yoon et al., 2018, Nguyen et al., 2019, Gordon et al., 2019, Chen et al., 2019b] and (ii) derive amortized conditional probability models [Garnelo et al., 2018, Gordon et al., 2019]. The underlying learning principle in both (i)-(ii) is to construct and maximize a predictive distribution (or conditional marginal likelihood) of the validation points given the training points, which, e.g., in supervised few-shot learning is written as $p_\theta(Y^v|X^v, X^t, Y^t) = \int p(Y^v|X^v, \psi_i)p_\theta(\psi_i|X^t, Y^t)d\psi_i = \frac{p_\theta(Y^v, Y^t|X^v, X^t)}{p_\theta(Y^t|X^t)}$. Here, $p_\theta(\psi_i|X^t, Y^t)$ is a posterior distribution over the task parameters ψ_i , after observing the training points, and θ is a meta parameter which for simplicity we assume to be

found by point estimation. However, this objective is difficult to approximate. Unlike the marginal likelihood on all task outputs $p_\theta(Y^v, Y^t|X^v, X^t)$, for which we can easily compute a lower bound, there is no tractable lower bound on the predictive conditional $p_\theta(Y^v|X^v, X^t, Y^t)$.⁴ This inherent difficulty with computing the predictive distribution has led to several approximations, i.e., methods of category (i) above, ranging from maximum a posteriori (MAP), Laplace, variational inference procedures (without bounds on the predictive conditional) and Stein variational gradient descent. The approaches of category (ii) try to directly model $p_\theta(Y^v|X^v, X^t, Y^t)$ without considering this as an approximation to an initial joint Bayesian model. Our VIB framework differs significantly from the predictive distribution principle, since VIB has an information theoretic motivation and it rigorously bounds an information bottleneck objective. VIB is also a fully tractable objective, thus avoiding the need to choose a particular approximate inference method and allowing us to rather focus on setting up the encoding procedure, as demonstrated with the GP example from Section 3.

Finally, regarding related works of GPs in meta learning, the ALPaCA method [Harrison et al., 2018] applied GPs to Bayesian linear regression, while Tossou et al. [2019] used kernel-based methods from a regularization rather than Bayesian perspective. Closer to our work, Patacchiola et al. [2020] and Snell and Zemel [2021] used GPs with deep neural kernels for few-shot classification. Our usage of GPs is different; e.g., our encoder amortization strategy can potentially deal with arbitrary likelihood functions and task output observations, while Patacchiola et al. [2020] assume a Gaussian likelihood for the binary class labels and Snell and Zemel [2021] consider the Pólya-gamma augmentation, which is tailored to classification problems.

5 EXPERIMENTS

Here we evaluate the algorithm from Section 3 (labeled “GP-VIB” in this section). To that end, we consider a standard set of meta-learning benchmarks: sinusoid regression and few-shot classification.

Sinusoid regression. We first evaluate the method on sinusoid regression, following the settings of Finn et al. [2017]. Each task involves regressing from the input to the output of a sine wave, where the amplitude and phase of the sinusoids vary across tasks.

We fit GP-VIB with regularization $\beta = 1$ and linear kernel. More in detail, the kernel feature vector $\phi(x; \theta)$ is obtained by the last hidden layer of the same architecture used in MAML [Finn et al., 2017]. Based on this M -dimensional feature vector $\phi(x; \theta)$, we obtain the kernel as $k_\theta(x, x') =$

⁴To obtain such a bound, we either need to have access to the intractable posterior $p_\theta(\psi_i|X^t, Y^t)$ or to upper bound the marginal likelihood on the training points $p_\theta(Y^t|X^t)$, which is hard.

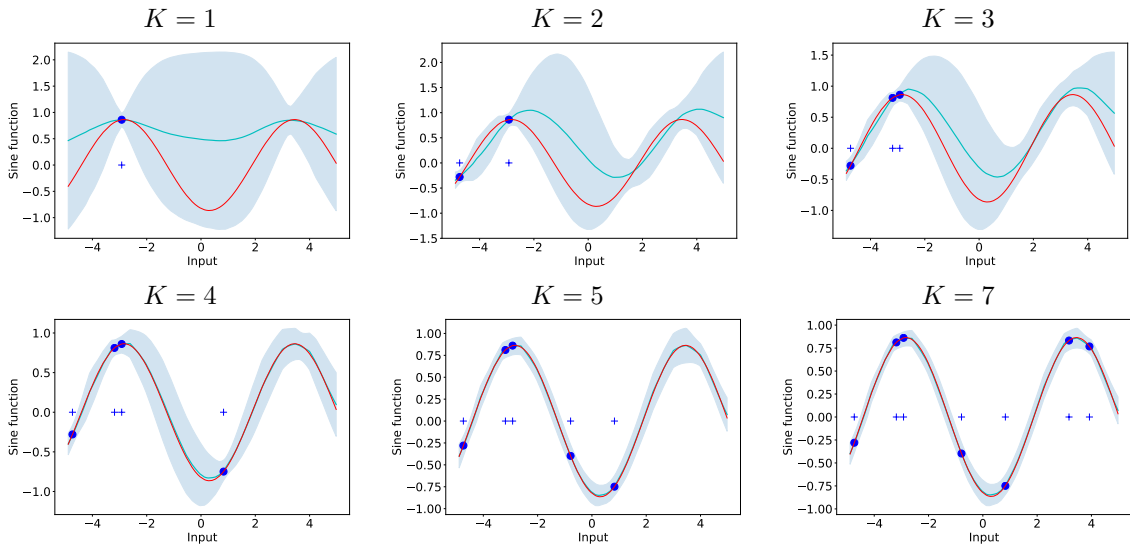


Figure 2: Comparison of the ground truth sinusoid function (red) with the mean prediction of GP-VIB (light blue). As the number of shots K increases (the observations are depicted by blue points), the uncertainty (shaded area) of the GP decreases and its mean predictions become closer to the ground truth. With just $K = 4$ points, the mean prediction of the GP closely matches the sinusoid function.

Method	$K = 5$	$K = 10$	$K = 20$
MAML-steps 1	0.600 ± 0.662	0.359 ± 0.015	0.228 ± 0.018
MAML-steps 5	0.311 ± 0.013	0.12 ± 0.006	0.06 ± 0.004
MAML-steps 10	0.280 ± 0.013	0.096 ± 0.005	0.043 ± 0.003
GP-VIB [ours]	0.02 ± 0.014	0.002 ± 0.001	0.001 ± 0.001

Table 1: MSE for K -shot sinusoid regression (mean values and 95% confidence intervals after 10 repeats). GP-VIB exhibits significantly smaller error than MAML.

$\frac{1}{M} \phi(x; \theta)^\top \phi(x'; \theta)$ (where the kernel variance σ_f^2 is fixed to $1/M$). Besides GP-VIB, we also fit MAML (with varying number of adaptation steps) for comparisons. In particular, following Finn et al. [2017], we meta train MAML with one gradient adaptation step, while at meta testing we consider several adaptation steps (1, 5, 10).

We report the performance of both methods in terms of K -shot mean squared error (MSE) in Table 1. We observe that GP-VIB significantly outperforms MAML, especially as K grows. Figure 2 explains this few-shot predictive ability of GP-VIB, as the GP posterior uncertainty decreases as K grows, and the posterior mean becomes very close to the ground truth sinusoid after only $K = 4$ shots.

Few-shot classification. We now evaluate GP-VIB on few-shot classification. To provide a fair comparison across different methods, we follow the unified training and evaluation protocol of Patacchiola et al. [2020] and implement GP-VIB building on their PyTorch code.⁵

⁵The code by Patacchiola et al. [2020] builds on the im-

We train both 1-shot and 5-shot versions of GP-VIB in four settings: Caltech-UCSD Birds (CUB) [Wah et al., 2011], mini-ImageNet, and two cross-domain transfer tasks—training on mini-ImageNet and testing on CUB, and training on Omniglot [Lake et al., 2011] and testing on EMNIST [Cohen et al., 2017]. The CUB dataset consists of 11788 images across 200 classes. We divide the dataset in 100 classes for training, 50 for validation, and 50 for testing [Chen et al., 2019a]. The mini-ImageNet dataset consists of a subset of 100 classes (600 images for each class) taken from the ImageNet dataset. We use 64 classes for training, 16 for validation and 20 for testing, as is common practice [Ravi and Larochelle, 2017, Chen et al., 2019a]. The Omniglot dataset contains 1623 characters taken from 50 different languages. Following the experimental protocol of Nichol et al. [2018], the number of classes is increased to 6492 by applying data augmentation and adding examples rotated by 90 degrees. We use 4114 classes for training. The EMNIST dataset contains single digits and characters from the English alphabet. The total 62 classes are divided into 31 for validation and 31 for test. The whole experimental protocol is taken from Patacchiola et al. [2020].

To fit GP-VIB, we employ the commonly used Conv4 architecture with 64 channels [Vinyals et al., 2016] and we use the Adam optimizer with the default learning rate of 10^{-3} . We use the simplified GP encoder described in Section 3.2 and we set the regularization coefficient $\beta = 0.001$ in all our experiments. We may obtain better results via hyperparameter search for each specific dataset, but using a single value

plementation of Chen et al. [2019a] and is available at <https://github.com/BayesWatch/deep-kernel-transfer>.

Method	CUB		mini-ImageNet	
	1-shot	5-shot	1-shot	5-shot
Feature Transfer	46.19 ± 0.64	68.40 ± 0.79	39.51 ± 0.23	60.51 ± 0.55
Baseline++ [Chen et al., 2019a]	61.75 ± 0.95	78.51 ± 0.59	47.15 ± 0.49	66.18 ± 0.18
MatchingNet [Vinyals et al., 2016]	60.19 ± 1.02	75.11 ± 0.35	48.25 ± 0.65	62.71 ± 0.44
ProtoNet [Snell et al., 2017]	52.52 ± 1.90	75.93 ± 0.46	44.19 ± 1.30	64.07 ± 0.65
MAML [Finn et al., 2017]	56.11 ± 0.69	74.84 ± 0.62	45.39 ± 0.49	61.58 ± 0.53
RelationNet [Sung et al., 2018]	62.52 ± 0.34	78.22 ± 0.07	48.76 ± 0.17	64.20 ± 0.28
DKT + Linear [Patacchiola et al., 2020]	60.23 ± 0.76	74.74 ± 0.22	48.44 ± 0.36	62.88 ± 0.46
DKT + CosSim [Patacchiola et al., 2020]	63.37 ± 0.19	77.73 ± 0.26	48.64 ± 0.45	62.85 ± 0.37
DKT + BNCosSim [Patacchiola et al., 2020]	62.96 ± 0.62	77.76 ± 0.62	49.73 ± 0.07	64.00 ± 0.09
GP-VIB + Linear [ours]	62.73 ± 0.59	77.78 ± 0.27	49.61 ± 0.29	65.84 ± 0.22
GP-VIB + CosSim [ours]	62.36 ± 0.77	78.35 ± 0.23	49.80 ± 0.52	65.60 ± 0.54

(a) Few-shot classification on CUB and mini-ImageNet.

Method	Omniglot → EMNIST		mini-ImageNet → CUB	
	1-shot	5-shot	1-shot	5-shot
Feature Transfer	64.22 ± 1.24	86.10 ± 0.84	32.77 ± 0.35	50.34 ± 0.27
Baseline++ [Chen et al., 2019a]	56.84 ± 0.91	80.01 ± 0.92	39.19 ± 0.12	57.31 ± 0.11
MatchingNet [Vinyals et al., 2016]	75.01 ± 2.09	87.41 ± 1.79	36.98 ± 0.06	50.72 ± 0.36
ProtoNet [Snell et al., 2017]	72.04 ± 0.82	87.22 ± 1.01	33.27 ± 1.09	52.16 ± 0.17
MAML [Finn et al., 2017]	72.68 ± 1.85	83.54 ± 1.79	34.01 ± 1.25	48.83 ± 0.62
RelationNet [Sung et al., 2018]	75.62 ± 1.00	87.84 ± 0.27	37.13 ± 0.20	51.76 ± 1.48
DKT + Linear [Patacchiola et al., 2020]	75.97 ± 0.70	89.51 ± 0.44	38.72 ± 0.42	54.20 ± 0.37
DKT + CosSim [Patacchiola et al., 2020]	73.06 ± 2.36	88.10 ± 0.78	40.22 ± 0.54	55.65 ± 0.05
DKT + BNCosSim [Patacchiola et al., 2020]	75.40 ± 1.10	90.30 ± 0.49	40.14 ± 0.18	56.40 ± 1.34
GP-VIB + Linear [ours]	76.01 ± 0.54	89.93 ± 0.23	39.52 ± 0.7	54.51 ± 0.45
GP-VIB + CosSim [ours]	74.58 ± 1.56	88.73 ± 0.75	40.70 ± 0.48	56.70 ± 0.62

(b) Cross-domain few-shot classification.

Table 2: Average accuracy and standard deviation (percentage, based on 3 independent runs) on few-shot classification (5-ways). The two methods with the best average accuracy are highlighted in bold. All results, except for GP-VIB, are taken from Patacchiola et al. [2020]. GP-VIB provides competitive accuracy, obtaining state-of-the-art results in some cases.

for all of them showcases the robustness of the method. The kernel feature vector $\phi(x; \theta)$ is obtained by the last hidden layer of the same architecture used in MAML. We consider two kernel functions: the standard linear kernel (labeled “GP-VIB + Linear”), $k_\theta(x, x') = \frac{\exp(v)}{M} \phi(x; \theta)^\top \phi(x'; \theta)$ and the cosine similarity kernel (“GP-VIB + CosSim”), $k_\theta(x, x') = \frac{\exp(v) \phi(x; \theta)^\top \phi(x'; \theta)}{\|\phi(x; \theta)\| \|\phi(x'; \theta)\|}$, where the out-scale parameter v is updated with a smaller learning rate (10^{-4}), exactly as in Patacchiola et al. [2020].

We compare GP-VIB against a series of baseline methods, including MAML as well as more recent methods. We evaluate each method in terms of the classification accuracy. We report the results in Table 2, together with the standard deviation obtained from three independent runs. We can observe that GP-VIB is a competitive method, and it exhibits state-of-the-art results in some cases. In the settings where GP-VIB is not the best performing method, its accuracy is close to the best one.

6 CONCLUSIONS

We have introduced an information theoretic framework for meta learning by using a variational approximation [Alemi et al., 2017, Chalk et al., 2016, Achille and Soatto, 2016] to the information bottleneck [Tishby et al., 1999]. Based on this VIB view, we have developed a memory-based meta learning method that uses GPs to obtain a non-parametric stochastic encoding representation. We have shown experimentally that this method outperforms MAML in few-shot sinusoid regression and provides competitive performance on four few-shot classification problems, where it gives state-of-the-art results in some cases.

While we have demonstrated our method in few-shot regression and classification, we believe that the scope of the information bottleneck for meta learning is much broader. For instance, a promising topic for future research is to consider applications in reinforcement learning.

Acknowledgements

We are grateful to Yutian Chen for his help with the implementation and for many insightful discussions and comments on the manuscript. We also thank Jakub Sygnowski for technical advice regarding the implementation, and Yee Whye Teh and Matthias Bauer for their useful comments.

References

- Alessandro Achille and Stefano Soatto. Information dropout: Learning optimal representations through noisy computation, 2016.
- Alex Alemi, Ian Fischer, Josh Dillon, and Kevin Murphy. Deep variational information bottleneck. In *ICLR*, 2017.
- David Barber and Felix Agakov. The IM algorithm: A variational approach to information maximization. In *Advances in Neural Information Processing Systems*, 2003.
- John Bronskill, Jonathan Gordon, James Requeima, Sebastian Nowozin, and Richard E. Turner. TaskNorm: Rethinking batch normalization for meta-learning. *arXiv preprint arXiv:2003.03284*, 2020.
- Christopher P. Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β -VAE. *arXiv preprint arXiv:1804.03599*, 2018.
- Matthew Chalk, Olivier Marre, and Gasper Tkacik. Relevant sparse codes with variational information bottleneck. In *Advances in Neural Information Processing Systems*, 2016.
- Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Frank Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *International Conference on Learning Representations*, 2019a.
- Yutian Chen, Abram L. Friesen, Feryal Behbahani, Arnaud Doucet, David Budden, Matthew W. Hoffman, and Nando de Freitas. Modular meta-learning with shrinkage. *arXiv preprint arXiv:1909.05557*, 2019b.
- Gregory Cohen, Saeed Afshar, Jonathan Tapon, and André van Schaik. Emnist: an extension of mnist to handwritten letters, 2017.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006. ISBN 0471241954.
- Bin Dai, Chen Zhu, Baining Guo, and David P. Wipf. Compressing neural networks using the variational information bottleneck. In *International Conference on Machine Learning*, 2018.
- Harrison Edwards and Amos Storkey. Towards a neural statistician. In *International Conference on Learning Representations*, 2017.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.
- Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, 2018.
- Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and S. M. Ali Eslami. Conditional neural processes. In *International Conference on Machine Learning*, 2018.
- Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard Turner. Meta-learning probabilistic inference for prediction. In *International Conference on Learning Representations*, 2019.
- Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical Bayes. In *International Conference on Learning Representations*, 2018.
- James Harrison, Apoorva Sharma, and Marco Pavone. Meta-learning priors for efficient online Bayesian regression. *arXiv preprint arXiv:1807.08912*, 2018.
- James Hensman, Max Zwiessle, and Neil D. Lawrence. Tilted variational Bayes. In *International Conference on Artificial Intelligence and Statistics*, 2014.
- Xu Hu, Pablo Moreno, Yang Xiao, Xi Shen, Guillaume Obozinski, and Neil Lawrence. Empirical Bayes transductive meta-learning with synthetic gradients. In *International Conference on Learning Representations*, 2020.
- Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Alexandre Lacoste, Boris Oreshkin, Wonchang Chung, Thomas Boquet, Negar Rostamzadeh, and David Krueger. Uncertainty in multitask transfer learning. *arXiv preprint arXiv:1806.07528*, 2019.
- Brenden M. Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua B. Tenenbaum. One shot learning of simple visual concepts. *Cognitive Science*, 33, 2011.
- Fei-Fei Li, Robert Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611, 2006.

- Cuong Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Uncertainty in model-agnostic meta-learning using variational inference. *arXiv preprint arXiv:1907.11864*, 2019.
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.
- Manfred Opper and Cédric Archambeau. The variational Gaussian approximation revisited. *Neural Comput.*, 21(3):786–792, March 2009.
- Massimiliano Patacchiola, Jack Turner, Elliot J. Crowley, Michael O'Boyle, and Amos J. Storkey. Bayesian meta-learning for the few-shot setting via deep kernels. In *Advances in Neural Information Processing Systems*, 2020.
- Carl Edward Rasmussen and Christopher KI Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Sachin Ravi and Alex Beaton. Amortized Bayesian meta-learning. In *International Conference on Learning Representations*, 2019.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Jake Snell and Richard Zemel. Bayesian few-shot classification with one-vs-each Pólya-Gamma augmented Gaussian processes. In *International Conference on Learning Representations*, 2021.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 2017.
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. In *Annual Allerton Conference on Communication, Control and Computing*, 1999.
- Prudencio Tossou, Basile Dura, Francois Laviolette, Mario Marchand, and Alexandre Lacoste. Adaptive deep kernel learning. *arXiv preprint arXiv:1905.12131*, 2019.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial Intelligence and Statistics*, 2016.
- Jaesik Yoon, Taesup Kim, Ousmane Dia, Sungwoong Kim, Yoshua Bengio, and Sungjin Ahn. Bayesian model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, 2018.