
Probabilistic task modelling for meta-learning

Cuong C. Nguyen¹

Thanh-Toan Do²

Gustavo Carneiro¹

¹Australian Institute for Machine Learning, University of Adelaide, Australia

²Department of Data Science and AI, Monash University, Australia

Abstract

We propose *probabilistic task modelling* – a generative probabilistic model for collections of tasks used in meta-learning. The proposed model combines variational auto-encoding and latent Dirichlet allocation to model each task as a mixture of Gaussian distribution in an embedding space. Such modelling provides an explicit representation of a task through its task-theme mixture. We present an efficient approximation inference technique based on variational inference method for empirical Bayes parameter estimation. We perform empirical evaluations to validate the *task uncertainty* and *task distance* produced by the proposed method through correlation diagrams of the prediction accuracy on testing tasks. We also carry out experiments of task selection in meta-learning to demonstrate how the task relatedness inferred from the proposed model help to facilitate meta-learning algorithms.

1 INTRODUCTION

The latest developments in machine learning have enabled the field to solve increasingly complex classification problems. Such complexity require high capacity models, which in turn need a massive amount of annotated data for training, resulting in an arduous, costly and even infeasible annotation process. This has, therefore, motivated the research of novel learning approaches, generally known as transfer learning, that exploit past experience (in the form of models learned from other training tasks) to quickly learn a new task using relatively small training sets.

Transfer-learning, and in particular, meta-learning, assumes the existent of a task environment where training and testing tasks are i.i.d. sampled from the same latent distribution. By modelling such environment through meta-parameters that are shared across all tasks, meta-learning can solve

an unseen task more accurately and efficiently by learning how to solve many tasks generated from the same distribution, even if each task contains a limited number of training examples. Meta-learning has, therefore, progressed steadily with many remarkable state-of-the-art results in several few-shot learning benchmarks [Vinyals et al., 2016, Snell et al., 2017, Finn et al., 2017, Yoon et al., 2018, Rusu et al., 2019, Allen et al., 2019]. However, current development of meta-learning focuses on solving tasks without providing understanding on how tasks are generated or correlated, potentially leading to sub-optimal solutions. In fact, there is a large variation of prediction performance made by various meta-learning algorithms reported in [Dhillon et al., 2020, Figure 1] or shown in Figure 1, suggesting that not all testing tasks are equally related to training tasks. This motivates our work to model and represent tasks in a latent “task-theme” space. The new task representation allows further downstream works, such as task similarity or active task selection, to be developed to gain insights into, or even improve, the prediction performance of different meta-learning algorithms.

In this paper, we propose **probabilistic task modelling** (PTM) – a graphical model that combines variational auto-encoding (VAE) [Kingma and Welling, 2014] and Gaussian latent Dirichlet allocation (LDA) [Das et al., 2015] – to **model tasks** used in meta-learning. Note that PTM itself is not a meta-learning method. With this modelling approach, the dataset associated with each task can be modelled as a mixture of finite Gaussian distributions, allowing to represent tasks in a latent “task-theme” simplex via its mixture coefficient vector. Such representation provides a convenient way to quantitatively measure “task uncertainty” or task relatedness, which can be utilised in active task selection for meta-learning.

2 RELATED WORK

The proposed approach is closely related to Task2Vec [Achille et al., 2019] when modelling tasks for meta-

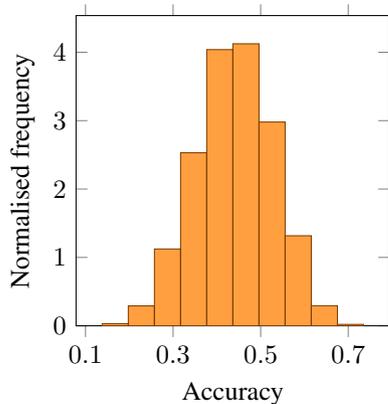


Figure 1: The results locally produced for MAML on 15,504 available 5-way 1-shot mini-ImageNet testing tasks vary from 20 to 70 percent accuracy, suggesting that not all testing tasks are equally related to training tasks.

learning. In Task2Vec, a task is represented by an embedding computed from the Fisher information matrix associated with the task-specific classifier. In PTM, a task is represented by the variational distribution of task-theme mixture, which is a part of the graphical model describing the task generation process. The two methods, therefore, differ at the modelling mechanism: Task2Vec follows a deterministic approach, while PTM is a probabilistic method. Such difference provides an advantage of PTM over Task2Vec, which includes modelling uncertainty into the task representation. In addition, PTM is more efficient than Task2Vec at inference when predicting task representation, since PTM only needs a single forward pass, while Task2Vec requires to re-train or fine-tune the task-specific classifier and calculate the Fisher information matrix for the task that needs to be presented.

Our work is related to task similarity estimation, which has been intensively studied in the field of multi-task learning. Some remarkable examples in this area include task-clustering using k-nearest neighbours [Thrun and O’Sullivan, 1996], or modelling common prior between tasks as a mixture of distributions [Bakker and Heskes, 2003, Xue et al., 2007]. Another approach is to formulate multi-task learning as a convex optimisation problem either to cluster tasks and utilise the clustering results to fast track the learning [Jacob et al., 2009], or to learn task relationship through task covariance matrices [Zhang and Yeung, 2012]. Other approaches provided theoretical guarantees when learning the similarity or relationship between tasks [Shui et al., 2019]. Recently, the *taskonomy* project [Zamir et al., 2018] was conducted to carry out extensive experiments on 26 computer-vision tasks to empirically analyse the correlation between those tasks. Other works [Tran et al., 2019, Nguyen et al., 2020b] take a slightly different approach by investigating the correlation of the label distributions

between the tasks of interest to measure task similarity. One commonality among all studies above is their reliance on task-specific classifiers which are used to quantify task relatedness. In contrast, our proposed method explicitly models tasks without the help of any task-specific classifier, making it more efficient in training and prediction.

Our work is also connected to finite mixture models [Pritchard et al., 2000], such as the *latent Dirichlet allocation* [Blei et al., 2003], which analyses and summarises text data in topic modelling, or categorises natural scenes in computer vision [Li and Perona, 2005]. LDA assumes that each document within a given corpus can be represented as a mixture of finite categorical distributions, where each categorical distribution is a latent topic shared across all documents. Training an LDA model or its variants on a large text corpus is challenging, so several approximate inference techniques have been proposed, ranging from mean-field variational inference (VI) [Blei et al., 2003], collapsed Gibbs’ sampling [Griffiths and Steyvers, 2004] and collapsed VI [Teh et al., 2007]. Furthermore, several online inference methods have been developed to increase the training efficiency for large corpora [Canini et al., 2009, Hoffman et al., 2010, Foulds et al., 2013]. Our work is slightly different from the modelling of the conventional LDA, where we do not use the data directly, but embed it into a latent space. In short, our approach is a combination of VAE [Kingma and Welling, 2014] and LDA to model the dataset associated with a task. Our approach considers “word” as continuous data, instead of the discrete data represented by a bag-of-word vector generally used by LDA-based topic modelling methods. The resultant model in the embedding latent space is, therefore, similar to the Gaussian LDA [Das et al., 2015] for word embedding in topic modelling.

3 PROBABILISTIC TASK MODELLING

To relate **task modelling** to **topic modelling**, we consider **a task** as **a document**, and **a data point** as **a word**. Given these analogies, we can use LDA [Blei et al., 2003] – a popular topic model – to model tasks for meta-learning. However, simply applying LDA for task modelling would not scale well with high-dimensional data and large datasets. We, therefore, propose to employ the VAE [Kingma and Welling, 2014] to reduce the dimension of the data, and use the inferred embeddings of data as words to model tasks. Due to the nature of VAE, the latent variables are often continuous, not discrete as the bag-of-words used in the conventional LDA. We, therefore, **replace the categorical word-topic distributions in LDA by Gaussian task-theme¹ distributions**. Given these assumptions, each task can be modelled as a mixture of K Gaussian task-themes, allowing to represent tasks by their inferred *task-theme mixture* vectors in the latent task-theme simplex as illustrated in Fig-

¹“Task-theme” is inspired by Li and Perona [2005]

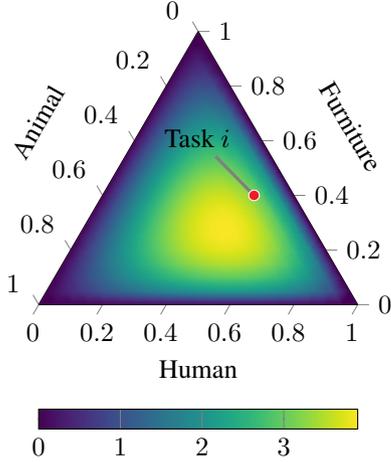


Figure 2: An example of a task-theme simplex where each task is represented by a 3-dimensional mixture vector.

ure 2. Hence, it is beneficial to utilise this representation for further downstream tasks, such as measuring task similarity.

The graphical model of the proposed PTM is shown in Figure 3, where there are T tasks, and each task consists of N data points, denoted as \mathbf{x} . To simplify the formulation and analysis, N is assumed to be fixed across all tasks, but the extension of varying N is trivial and can be implemented straightforwardly. Under these assumptions, a task can be generated as follows:

- Initialise the Dirichlet prior for task-theme mixture: $\{\alpha_k\}_{k=1}^K$, where $\alpha \in \mathbb{R}_+$
- Initialise means and covariance matrices of K Gaussian task-themes $\{\mu_k, \Sigma_k\}_{k=1}^K$, where $\mu_k \in \mathbb{R}^D$, $\Sigma_k \in \mathbb{R}^{D \times D}$ is positive definite matrix, and D is the dimension of the data embedding
- For task \mathcal{T}_i in the collection of T tasks:
 - Choose a task-theme mixture vector: $\pi_i \sim \text{Dirichlet}(\pi; \alpha)$
 - For data point n -th of task \mathcal{T}_i :
 - * Choose an task-theme assignment one-hot vector: $\mathbf{z}_{in} \sim \text{Categorical}(\mathbf{z}; \pi_i)$
 - * Choose an embedding of the data point: $\mathbf{u}_{in} \sim \mathcal{N}(\mathbf{u}; \mu_k, \Sigma_k)$, where: $z_{ink} = 1$
 - * Generate the data point from a decoder h parameterised by θ : $\mathbf{x}_{in} = h(\mathbf{u}_{in}; \theta)$.

To model tasks according to the task generation described above, we need to infer the task-agnostic (or meta) parameters of interest, namely μ, Σ, α and θ . However, due to the complexity of the graphical model shown in Figure 3, the exact inference for the posterior $p(\mu, \Sigma, \alpha, \theta | \mathbf{x})$ is intractable, and therefore, the estimation must rely on approximate inference. For simplicity, maximum likelihood estimation

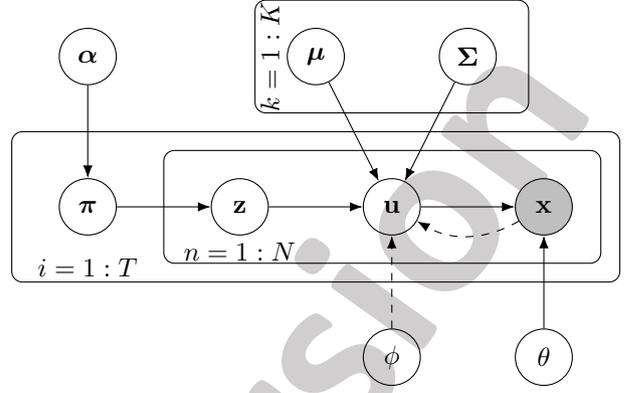


Figure 3: The graphical model used in task modelling. The solid arrows denote data generation, while the dashed arrows stand for inference. The boxes are “plates” representing replicates. The shading nodes denote observable variables, while the white nodes denote latent variables.

(MLE) is used as the objective function:

$$\max_{\mu, \Sigma, \alpha, \theta} \ln p(\mathbf{x} | \mu, \Sigma, \alpha, \theta). \quad (1)$$

Although MLE simplifies the learning for the meta-parameters of interest, the log-likelihood in (1) is still difficult to evaluate for the optimisation. One workaround solution is to find its lower-bound, and maximise the lower-bound instead of maximising the log-likelihood itself. This approach is analogous to the variational inference, which has been widely used to infer the latent parameters of VAE and LDA models.

Since the proposed PTM is a combination of VAE and LDA, the derivation for the lower-bound of the likelihood in (1) can be divided into 2 steps, where the first step is analogous to the lower bound of a VAE, and the second step is similar to the plain LDA model.

In the first step, the latent variable \mathbf{u} is introduced, so that the log-likelihood $\ln p(\mathbf{x} | \mu, \Sigma, \alpha, \theta)$ can be bounded below by Jensen’s inequality:

$$\ln p(\mathbf{x} | \mu, \Sigma, \alpha, \theta) \geq L_{\text{VAE}}, \quad (2)$$

where the lower-bound is defined as:

$$L_{\text{VAE}} = \mathbb{E}_{q(\mathbf{u})} [\ln p(\mathbf{x} | \mathbf{u}, \theta) + \ln p(\mathbf{u} | \mu, \Sigma, \alpha) - \ln q(\mathbf{u})], \quad (3)$$

with $q(\mathbf{u})$ being the variational distribution for the latent variable \mathbf{u} .

Following the conventional VI for VAE [Kingma and Welling, 2014], the variational distribution for the data embedding \mathbf{u} is assumed to be a Gaussian distribution with diagonal covariance matrix:

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{m}, (\mathbf{s})^2 \mathbf{I}), \quad (4)$$

where \mathbf{I} represents an identity matrix.

In addition, the parameters \mathbf{m} and \mathbf{s} , which represent the distribution encoding the data \mathbf{x} , are modelled by a neural network (also known as an encoder) f parameterised by ϕ :

$$[\mathbf{m} \ \mathbf{s}]^\top = f(\mathbf{x}; \phi). \quad (5)$$

Hence, instead of maximising the marginal log-likelihood in (1), the lower-bound in (3) is maximised, resulted in the alternative objective:

$$\max_{\mu, \Sigma, \alpha, \theta} \max_{\phi} L_{\text{VAE}}. \quad (6)$$

One difficulty in maximising the lower-bound in (6) is the evaluation for the embedding prior $\ln p(\mathbf{u}|\mu, \Sigma, \alpha)$ in Eq. (3). In vanilla VAE, the embedding prior is often modelled as some standard distributions, such as Gaussian or Beta, resulting in a tractable solution. In this paper, the prior is modelled as a Gaussian mixture model, making the solution intractable. However, since this prior is the marginal log-likelihood in the conventional LDA model, we can apply techniques developed for LDA methods to approximate this term. Here, we employ the VI approach in which the term is bounded below by Jensen’s inequality:

$$\ln p(\mathbf{u}|\mu, \Sigma, \alpha) \geq L_{\text{LDA}}(\mathbf{u}, q(\mathbf{z}, \pi)), \quad (7)$$

where:

$$L_{\text{LDA}}(\mathbf{u}, q(\mathbf{z}, \pi)) = \mathbb{E}_{q(\mathbf{z}, \pi)} [\ln p(\mathbf{u}|\mathbf{z}, \mu, \Sigma) + \ln p(\mathbf{z}|\pi) + \ln p(\pi|\alpha) - \ln q(\mathbf{z}) - \ln q(\pi)], \quad (8)$$

with $q(\mathbf{z}, \pi)$ being the variational distribution for \mathbf{z} and π . This corresponds to the second step in the derivation.

Similar to LDA [Blei et al., 2003], the variational distribution $q(\mathbf{z}, \pi)$ is assumed to be fully factorised and followed the conjugate priors:

$$q(\mathbf{z}, \pi) = \prod_{i=1}^T q(\pi_i; \gamma_i) \prod_{n=1}^N q(\mathbf{z}_{in}; \mathbf{r}_{in}), \quad (9)$$

where:

$$q(\pi_i; \gamma_i) = \text{Dirichlet}(\pi_i; \gamma_i) \quad (10)$$

$$q(\mathbf{z}_{in}; \mathbf{r}_{in}) = \text{Categorical}(\mathbf{z}_{in}; \mathbf{r}_{in}), \quad (11)$$

with \mathbf{r} and γ being the parameters of the variational distribution $q(\mathbf{z}, \pi)$.

In practice, $q(\mathbf{z}, \pi)$ is obtained as the maximiser of the lower-bound $L_{\text{LDA}}(\mathbf{u}, q(\mathbf{z}, \pi))$ on the embedding data \mathbf{u} . It is, however, inapplicable in this case, since the data embedding \mathbf{u} is used twice: one to optimise $q(\mathbf{z}, \pi)$, and the other is to optimise the objective in (6), which may result in overfitting. To avoid this issue, we employ the *empirical*

Bayes approach relying on the train-test split method, where one half of data in a task, denoted as $\mathbf{u}^{(t)}$, is used to obtain $q(\mathbf{z}, \pi)$, while the other half, denoted as $\mathbf{u}^{(v)}$, is used for the optimisation in (6). This approach is analogous to the *empirical Bayes* meta-learning [Finn et al., 2017, Nguyen et al., 2020a], where one part of data is used for task-adaptation (often known as “inner-loop”), while the other part is used to learn the meta-parameter (often known as “outer-loop”).

Given this modelling approach, the objective function can be formally written as a bi-level optimisation:

$$\max_{\mu, \Sigma, \alpha, \theta, \phi} L(\mathbf{u}^{(v)}, q^*(\mathbf{z}, \pi))$$

subject to:

$$q^*(\mathbf{z}, \pi) = \arg \max_{q(\mathbf{z}, \pi)} \mathbb{E}_{q(\mathbf{u}^{(t)}; \phi)} \left[L_{\text{LDA}}(\mathbf{u}^{(t)}, q(\mathbf{z}, \pi)) \right], \quad (12)$$

where

$$L(\mathbf{u}^{(v)}, q^*(\mathbf{z}, \pi)) = \mathbb{E}_{q(\mathbf{u}^{(v)}; \phi)} \left[L_{\text{LDA}}(\mathbf{u}^{(v)}, q^*(\mathbf{z}, \pi)) + \ln p(\mathbf{x}^{(v)}|\mathbf{u}^{(v)}, \theta) - \ln q(\mathbf{u}^{(v)}; \phi) \right]. \quad (13)$$

Due to the assumptions made in Eqs. (4), (10) and (11), prior conjugate can be applied to simplify the evaluation for all the terms in (8) w.r.t. the variational distribution $q(\cdot)$. Details of the evaluation can be referred to Appendix A. In addition, the optimisation for the meta-parameters in (12) is based on gradient ascent, and carried out in two steps, resulting in a process analogous to the expectation-maximisation (EM) algorithm. In the E-step (corresponding to the optimisation for the lower-level in (12)), the task-specific variational-parameters \mathbf{r} and γ are iteratively updated, while holding the meta-parameters $\mu, \Sigma, \alpha, \theta$ and ϕ fixed. In the M-step (corresponding to the optimisation for the upper-level), the meta-parameters are updated using the values of the task-specific variational-parameters obtained in the E-step. Note that the inference for the task-theme parameters μ and Σ are similar to the estimation of Gaussian mixture model [Bishop, 2006, Chapter 9]. Please refer to Appendix B for more details on the optimisation.

Conventionally, the iterative updates in the E-step and M-step require a full pass through the entire collection of tasks. This is, however, very slow and even infeasible since T is often in the magnitude of millions. We, therefore, propose an online VI inspired by the online learning for LDA [Hoffman et al., 2010] to infer the meta-parameters. For each task \mathcal{T}_i , we perform the EM to obtain the “task-specific” parameters (denoted by a tilde on top of variables) that are locally optimal for that task. The “meta” parameters of interest are then updated as a weighted average between their previous

values and the ‘‘task-specific’’ values:

$$\begin{aligned}\boldsymbol{\mu} &\leftarrow (1 - \rho_i)\boldsymbol{\mu} + \rho_i\tilde{\boldsymbol{\mu}} \\ \boldsymbol{\Sigma} &\leftarrow (1 - \rho_i)\boldsymbol{\Sigma} + \rho_i\tilde{\boldsymbol{\Sigma}} \\ \boldsymbol{\alpha} &\leftarrow \boldsymbol{\alpha} - \rho_i \underbrace{\tilde{\boldsymbol{\alpha}}_i}_{\mathbf{H}^{-1}\mathbf{g}},\end{aligned}\quad (14)$$

where $\rho_i = (\tau_0 + i)^{-\tau_1}$ with $\tau_0 \geq 0$ and $\tau_1 \in (0.5, 1]$ [Hoffman et al., 2010], \mathbf{g} is the gradient of \mathcal{L}_{LDA} w.r.t. $\boldsymbol{\alpha}$, and \mathbf{H} is the Hessian matrix. The learning for the encoder ϕ and the decoder θ follows the conventional learning by stochastic gradient ascent. The complete learning algorithm for the proposed probabilistic task modelling is shown in Algorithm 1.

Also, instead of updating the meta-parameters as in (14) when observing a single task, we use multiple or a mini-batch of tasks to reduce the effect of measurement noise. The mini-batch version requires a slight modification in the formulation presented above, where we calculate the average of all ‘‘task-specific’’ parameters for tasks in the same mini-batch, and use that as the ‘‘task-specific’’ value to update the corresponding ‘‘meta’’ parameters.

Although the ‘‘reconstruction’’ term $\ln p(\mathbf{x}^{(v)}|\mathbf{u}^{(v)}, \theta)$ in (12) is used to model the likelihood of un-labelled data, it can straightforwardly be extended to a labelled data pair $\{\mathbf{x}^{(v)}, \mathbf{y}^{(v)}\}$ by introducing the parameter \mathbf{w} of a classifier. In that case, the ‘‘reconstruction’’ term can be expressed as:

$$\begin{aligned}\ln p(\mathbf{x}^{(v)}, \mathbf{y}^{(v)}|\mathbf{u}^{(v)}, \theta, \mathbf{w}) \\ = \underbrace{\ln p(\mathbf{y}^{(v)}|\mathbf{u}^{(v)}, \mathbf{w})}_{\text{negative classification loss}} + \underbrace{\ln p(\mathbf{x}^{(v)}|\mathbf{u}^{(v)}, \theta)}_{\text{negative reconstruction loss}}.\end{aligned}\quad (15)$$

In general, \mathbf{w} can be either a task-specific parameter generated from an additional meta-parameter shared across all tasks – corresponding to *empirical Bayes* meta-learning (e.g. using train-test split to learn hyper-parameters) algorithms [Finn et al., 2017, Nguyen et al., 2020a], or a meta-parameter itself – corresponding to *metric* meta-learning [Vinyals et al., 2016, Snell et al., 2017]. For simplicity, we will use the latter approach relying on the prototypical network [Snell et al., 2017] with Euclidean distance on the data embedding \mathbf{u} , to calculate the classification loss on labelled data. This reduces the need to introduce an additional parameter \mathbf{w} into our modelling.

TASK REPRESENTATION

Given the inferred meta-parameters, including the task-themes $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K$, the Dirichlet prior $\{\boldsymbol{\alpha}_l\}_{l=1}^L$, the encoder ϕ and the decoder θ , we can embed the data of a task into a latent space, and calculate its variational Dirichlet posterior of the task-theme *mixing coefficients* $q(\boldsymbol{\pi}; \boldsymbol{\gamma}_i)$. The obtained distribution can be used represent the corresponding task in the latent task-theme simplex as illustrated in

Algorithm 1 Online probabilistic task modelling

```

1: procedure TRAINING
2:   Initialise LDA parameters:  $\{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \boldsymbol{\alpha}_k\}_{k=1}^K$ 
3:   Initialise encoder  $\phi$  and decoder  $\theta$ 
4:   for each mini-batch of  $T_{\text{mini}}$  tasks do
5:     for  $i = 1 : T_{\text{mini}}$  do
6:       Split data into  $\{\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)}\}$  and  $\{\mathbf{x}_i^{(v)}, \mathbf{y}_i^{(v)}\}$ 
7:        $\mathbf{m}_i^{(t)}, \mathbf{s}_i^{(t)} \leftarrow f(\mathbf{x}_i^{(t)}; \phi)$ 
8:        $\mathbf{m}_i^{(v)}, \mathbf{s}_i^{(v)} \leftarrow f(\mathbf{x}_i^{(v)}; \phi)$ 
9:        $\boldsymbol{\gamma}, \mathbf{r} \leftarrow \text{E-STEP}(\mathcal{N}(\mathbf{u}; \mathbf{m}_i^{(t)}, (\mathbf{s}_i^{(t)})^2\mathbf{I}))$ 
10:      Calculate  $\mathcal{L}(\mathbf{u}_i^{(v)}, q_i^*(\mathbf{z}, \boldsymbol{\pi})) \triangleright \text{Eq. (13)}$ 
11:      Calculate ‘‘local’’ task-themes  $\tilde{\boldsymbol{\mu}}_i, \tilde{\boldsymbol{\Sigma}}_i, \tilde{\boldsymbol{\alpha}}_i$ 
12:    end for
13:     $\bar{\mathcal{L}} = \frac{1}{T} \sum_{i=1}^T \mathcal{L}(\mathbf{u}_i^{(v)}, q_i^*(\mathbf{z}, \boldsymbol{\pi}))$ 
14:     $\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha} \leftarrow \text{online\_LDA}(\tilde{\boldsymbol{\mu}}_{1:T}, \tilde{\boldsymbol{\Sigma}}_{1:T}, \tilde{\boldsymbol{\alpha}}_{1:T})$ 
15:     $\theta, \phi \leftarrow \text{SGD}(-\bar{\mathcal{L}}) \triangleright \text{gradient ascent}$ 
16:  end for
17:  return  $\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \theta, \phi$ 
18: end procedure

19: procedure E-STEP( $\mathcal{N}(\mathbf{u}; \mathbf{m}, \mathbf{s}^2\mathbf{I})$ )
20:   Initialise  $\mathbf{r}, \boldsymbol{\gamma}$ 
21:   repeat
22:     calculate the un-normalised  $r_{ink} \triangleright \text{Eq. (25)}$ 
23:     normalise  $\mathbf{r}_{in}$  such that  $\sum_{k=1}^K r_{ink} = 1$ 
24:     calculate  $\gamma_{ik} \triangleright \text{Eq. (29)}$ 
25:   until  $\frac{1}{K} |\text{change in } \boldsymbol{\gamma}| < \text{threshold}$ 
26:   return  $\boldsymbol{\gamma}, \mathbf{r}$ 
27: end procedure

```

Figure 2. This new representation of tasks has two advantages comparing to the recently proposed task representation Task2Vec [Achille et al., 2019]: (i) it explicitly models and represents tasks without the need of any pre-trained networks to use as a ‘‘probe’’ network, and (ii) it uses a probability distribution, instead of a vector as in Task2Vec, allowing to include modelling uncertainty when representing tasks. Given the probabilistic nature of PTM, we can use the entropy of the inferred task-theme mixture distribution $q(\boldsymbol{\pi}; \boldsymbol{\gamma}_i)$ as a measure of *task uncertainty*. In section 4.1, we empirically show that this measure correlates to the generalisation or test performance.

In addition, the representation produced by PTM can be used to quantitatively analyse the similarity or *distance* between two tasks i and j through a divergence between $q(\boldsymbol{\pi}; \boldsymbol{\gamma}_i)$ and $q(\boldsymbol{\pi}; \boldsymbol{\gamma}_j)$. Commonly, symmetric distances, such as Jensen-Shannon divergence, Hellinger distance, or earth’s mover distance are employed to calculate the divergence between distributions. However, it is argued that similarity should be represented as an asymmetric measure [Tversky, 1977]. This is reasonable in the context of

transfer learning, since knowledge gained from learning a difficult task might significantly facilitate the learning of an easy task, but the reverse might not always have the same level of effectiveness. In light of asymmetric distance, we decide to use Kullback-Leibler (KL) divergence, denoted as $D_{\text{KL}}[\cdot|\cdot]$, to measure *task distance*. As $D_{\text{KL}}[P||Q]$ is defined as the information lost when using a code optimised for Q to encode the samples of P , we, therefore, calculate $D_{\text{KL}}[q(\boldsymbol{\pi}; \gamma_{T+1})||q(\boldsymbol{\pi}; \gamma_i)]$, where $i \in \{1, \dots, T\}$, to assess how the training task \mathcal{T}_i differs from the learning of the novel task \mathcal{T}_{T+1} .

4 EXPERIMENTS

In this section, we empirically validate the two properties of PTM – *task uncertainty* and *task distance* – through task distance matrix and correlation diagrams. We also show two applications of the proposed approach used in active task selection for inductive and transductive life-long meta-learning. The experiments are based on the n -way k -shot tasks formed from Omniglot [Lake et al., 2015] and mini-ImageNet [Vinyals et al., 2016] – the two widely used datasets to evaluate the performance of meta-learning algorithms.

The Omniglot dataset consists of 1623 different handwritten characters from 50 different alphabets, where each character was drawn in black and white by 20 different people. Instead of using random train-test split that mixes all characters, the original split [Lake et al., 2015] is used to yield finer-grained classification tasks. In addition to the task forming based on randomly mixing characters of many alphabets, the two-level hierarchy of alphabets and characters are utilised to increase the difficulty of the character classification. Note that no data augmentations, such as rotating images by multiples of 90 degrees, is used throughout the experiments. Also, all images are down-sampled to 64-by-64 pixel² to simplify the image reconstruct in the decoder.

The mini-ImageNet dataset comprises a small version of ImageNet, which contains 100 classes taken from ImageNet, and each class has 600 colour images. We follow the common train-test split that uses 64 classes for training, 16 classes for validation, and 20 classes for testing [Ravi and Larochelle, 2017]. Similar to Omniglot, all images are also in 64-by-64 pixel².

The encoder used in the experiments consists of 4 convolutional modules, where each module has a convolutional layer with 4-by-4 filters and 2-by-2 stride, followed by a batch normalisation and a leaky rectified linear activation function with a slope of 0.01. The output of the last convolutional layer is flattened and connected to a fully connected layer to output the desired dimension for the latent variable \mathbf{u} . The decoder is designed similarly, except that the convolutional operator is replaced by the corresponding transposed

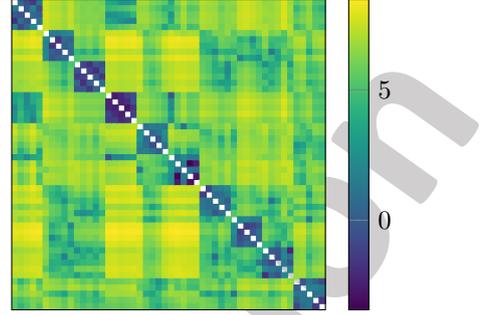


Figure 4: The matrix of log KL distances between Omniglot tasks shows that tasks that are generated from the same alphabet are closer together, denoted as the dark green blocks along the diagonal. The matrix is asymmetric due to the asymmetry of the KL divergence used as the task distance.

convolution. For the Omniglot dataset, the number of filters within each convolutional layer of the encoder is 8, 16, 32, and 64, respectively, and the dimension of \mathbf{u} is 64. For mini-ImageNet dataset, these numbers are 32, 64, 128 and 256, and the dimension of \mathbf{u} is 128. The reconstruction loss follows the negative log-likelihood of the continuous Bernoulli distribution [Loaiza-Ganem and Cunningham, 2019], which is often known as binary cross-entropy, while the classification loss is based on the prototypical network used in metric learning. The training subset of each task, $\mathbf{u}_i^{(t)}$, is used to calculate the class prototypes, and the classification loss is based on the soft-max function of the distances between the encoding of each input image to those prototypes [Snell et al., 2017]. The optimiser used is Adam with the step size of 2×10^{-4} to optimise the parameters of the encoder and decoder after every mini-batch consisting of 20 tasks. For the LDA part, a total of $K = 8$ task-themes is used. The Dirichlet prior is assumed to be symmetric with a concentration $\alpha = 1.1$ across both datasets². The parameters of the learning rate used in the online LDA are $\rho_0 = 10^6$ and $\rho_1 = 0.5$. A total of 10^6 episodes are used to train PTM on both datasets. We note that setting $\alpha > 1$ enforces every task to be modelled as a mixture of many task-themes, avoiding the task-themes collapsing into a single task-theme during training. The phenomenon of task-theme collapse when $\alpha < 1$ is not observed in LDA, but in PTM due to the integration of VAE. At the beginning of training, the encoder is inadequate, producing mediocre embedding features. The resulting features, combined with $\alpha < 1$, makes a task more likely to be represented by a single task-theme. By learning solely from that task-theme, the encoder is pushed to bias further toward to that task-theme, making only one task-theme distribution updated, while leaving others unchanged. When $\alpha > 1$, all the task-themes contribute to the representation of a task, so they can be learnt along with the encoder.

²https://github.com/cnguyen10/probabilistic_task_modelling

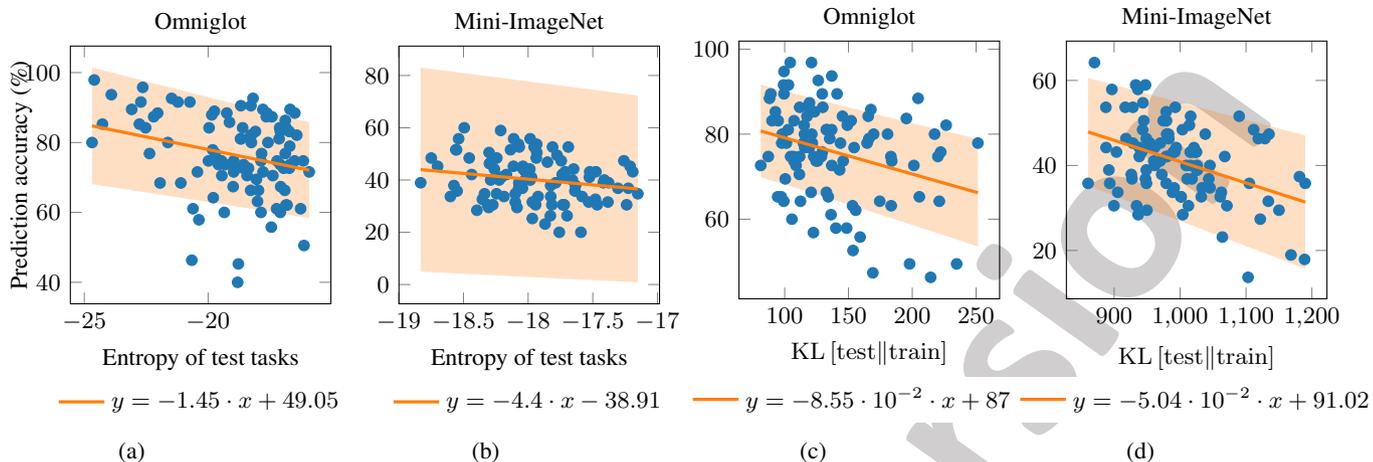


Figure 5: Correlation diagrams between prediction accuracy made by MAML on 100 5-way 1-shot testing tasks versus: (a) and (b) entropy of the inferred task-theme mixture distributions, and (c) and (d) the KL distances from testing to training tasks. The results show that largest the task entropy or distances, the worse the testing performance. The blue dots are the prediction made the MAML and PTM, the solid line is the mean of Bayesian Ridge regression, and the shaded areas correspond to ± 1 standard deviation around the mean.

4.1 TASK DISTANCE MATRIX AND CORRELATION DIAGRAMS

Task distance matrix is used as one of the tools to qualitatively validate the prediction made by PTM. In particular, the hypothesis is that the PTM would predict small distances for tasks that are close together. Since the “labels” specifying the closedness of tasks are unknown, we utilise the hierarchical structure of Omniglot dataset to form tasks. Each task is generated by firstly sampling an alphabet, and then choosing characters in that alphabet. Under this strategy, tasks formed from the same alphabet would have small distances comparing to tasks from different alphabets. Figure 4 shows the task distances between 50 testing tasks of Omniglot dataset, where each block of 5 tasks on rows and columns of the task distance matrix corresponds to a group of tasks sampled from the same alphabet. The result, especially the square 5-task-by-5-task blocks along the diagonal, agrees well with the hypothesis. Note that the distance matrix shown in Figure 4 is asymmetric due to the asymmetric nature of the KL divergence used to measure task distance.

We use a correlation diagram between prediction accuracy made by MAML and the *task entropy* produced by PTM as another verification. Since the *task entropy* denotes the uncertainty when modelling a task, we hypothesise that it proportionally relates to the difficulty when learning that task. To construct the correlation diagram, we firstly train a meta-learning model based on MAML using the training tasks of the two datasets, and evaluating the performance on 100 random testing tasks. Secondly, we calculate the *task entropy* for those 100 testing tasks. Finally, we plot the prediction accuracy and task entropy in Figures 5a and 5b. The results on both datasets show that the higher the task un-

certainly, the worse the test performance. This observation, therefore, agrees with our hypothesis about *task entropy*.

We conduct another correlation diagram between training-testing task distance and the test performance to verify further the proposed PTM. Our hypothesis is the inverse proportion between training-testing task distance and prediction accuracy. A similar experiment as in task uncertainty is carried out with a modification in which the task uncertainty is replaced by the average KL divergence between all training tasks to each testing task. Due to the extremely large number of training tasks, e.g. more than 10^{12} unique 5-way tasks can be generated from both the two datasets, the calculation of the distance measure is infeasible. To make the training and testing tasks manageable, we randomly generate 10,000 tasks for training, and 100 tasks for testing. This results in 1,000,000 distances, which can be calculated in parallel with multiple computers. A testing task can be represented in the correlation diagram through its prediction accuracy and the average KL distance to training tasks, which is defined as:

$$\bar{D}_{\text{KL}}(\gamma_{T+1}) = \frac{1}{T} \sum_{i=1}^T D_{\text{KL}}[q(\boldsymbol{\pi}; \gamma_{T+1}) \| q(\boldsymbol{\pi}; \gamma_i)].$$

The correlation diagrams for both datasets are then plotted in Figures 5c and 5d. The results agree well with our hypothesis, in which the further a testing task is from the training tasks, the worse the prediction accuracy. This enables us to use the new representation produced by PTM to analyse task similarity.

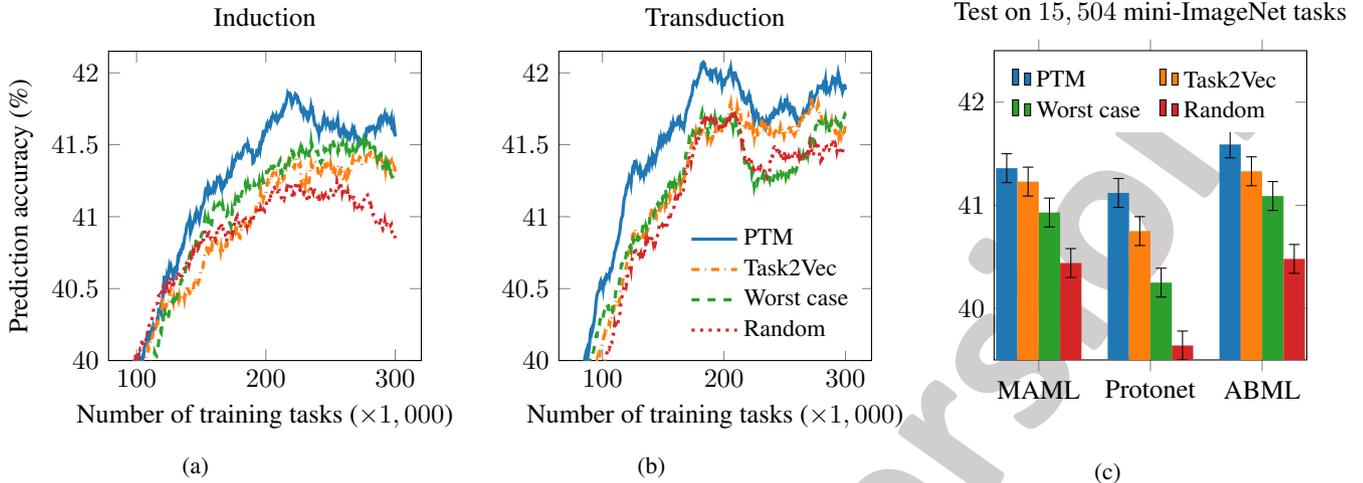


Figure 6: Exponential weighted moving average (EWMA) of prediction accuracy made by MAML following the lifelong learning for 100 random 5-way 1-shot tasks sampled from mini-ImageNet testing set: (a) inductive setting, and (b) transductive setting. The EWMA weight is set to 0.98 to smooth the noisy signal. (c) Prediction accuracy made by models trained on different task selection approaches on all 5-way 1-shot testing tasks generated from mini-ImageNet. The error bars correspond to 95 percent confident interval.

4.2 LIFELONG FEW-SHOT META-LEARNING

To further evaluate PTM, we conduct experiments following the lifelong learning framework [Ruvolo and Eaton, 2013] with slight modification where the supervised tasks are replaced by 5-way 1-shot learning episodes. More precisely, the setting consists of a meta-learning model and a pool of T_{pool} tasks. At each time step, a task selected from the pool is used to update the meta-learning model, and discarded from the pool. A new task is then added to the pool to maintain T_{pool} tasks available for learning. The criterion for selecting a task to update the meta-learning model will depend on the objective of interest. Two common objectives often observed in practice are:

- **Induction:** the selected training task is expected to encourage the meta-learning model to be able to rapidly adapt to **any** future task,
- **Transduction:** the selected training task is targeted toward one or many **specific** testing tasks.

In the induction setting, the performance of the meta-learning model trained on tasks selected by PTM is compared with three baselines: Task2Vec [Achille et al., 2019], the “worst-case” approach [Collins et al., 2020] and random selection. For the PTM, the selection criteria is based on the task entropy specified in section 3, where the training task with highest entropy is chosen for the learning. For Task2Vec, tasks with large embedding norm are reported as difficult to learn. Hence, we pick the one with the largest L1 norm produced by Task2Vec as the training task. Originally, Task2Vec requires fine-tuning a pre-trained network (known as probe network) on labelled data of a task. This

fine-tuning step is, however, infeasible for few-shot learning due to the insufficient number of labelled data. We address this issue by training a MAML-based network to use as a probe network. When given few-shot data of a training task, the MAML-based probe network perform gradient update to adapt to that task. The task-specific embedding can, therefore, be calculated using the adapted probe network. We follow the Monte Carlo approach specified in the public code of Task2Vec to calculate the corresponding task embedding. For the “worst-case” approach, the training task that results in the highest loss for the current meta-learning model is selected. Due to this nature, the “worst-case” approach requires to evaluate all losses for each task in the pool at every time step, leading to an extensive computation and might not scale well when the number of tasks in the pool is large. For simplicity, we use MAML to train the meta-learning model of interest for each selection strategy.

The transduction setting follows a similar setup as the induction case, but the testing tasks, including the labelled and unlabelled data, are known during training. For PTM, the average KL distances between all testing tasks to each training task in the task pool are calculated, and the training task with smallest average distance is selected. For Task2Vec, the proposed cosine distance between normalised task embeddings is used to calculate the average distance between all testing tasks to each training task [Achille et al., 2019]. Similar to PTM, the training task with the smallest distance is prioritised for the learning. For the “worst-case” approach, the entropy of the prediction \hat{y} on C -way testing tasks is

used as the measure:

$$S_{T+1} = - \sum_{c=1}^C \hat{y}_c \ln \hat{y}_c,$$

and the task that contributed to the highest entropy at prediction is chosen [MacKay, 1992]. The “worst-case” approach, therefore, requires T_{pool} trials at every time step. In each trial, the current meta-model is adapted to each training task in the pool, and then the average prediction entropy on all testing tasks is calculated. This results in an extremely extensive computation.

Four MAML-based meta-learning models are initialised identically and trained on the tasks selected from a pool of $T_{\text{pool}} = 200$ tasks according to the four criteria mentioned above. Figures 6a and 6b show the testing results on 100 random mini-ImageNet tasks after every 500 time steps. Note that the plotted results are smoothed by the exponential weighted moving average with a weight of 0.98 to ease the visualisation. In general, PTM, Task2Vec and “worst-case” can generalise better than random task selection. In addition, the model trained with tasks chosen by PTM performs slightly better than Task2Vec and the “worst-case” approach in both settings. This observation might be explained based on the designated purpose of Task2Vec and the “worst-case” approach. Task2Vec requires a sufficient number of labelled data to fine-tune its probe network to calculate task embedding. Hence, it might not work well in few-shot learning. For the “worst-case”, tasks are selected according to a measure based on the current meta-model without taking task relatedness into account. PTM, however, has a weakness in active selection since the approach only focuses on task uncertainty or task similarity without considering the current state of the meta-learning model. Nevertheless, PTM still provides a good selection criterion comparing to Task2Vec and the “worst-case” approaches. Note that although the active task selection is able to select the best task within the pool, there might be the case where all remaining tasks in the pool are uninformative, resulting in overfitting as observed in Figure 6a. However, for simplicity, no additional mechanism is integrated to decide whether to learn from the selected task, or simply discarded from the pool. We believe that adding L2 regularisation or applying early stopping based on a validation set of tasks will help with this overfitting issue.

To further compare, we implement two additional meta-learning algorithms: Prototypical Networks [Snell et al., 2017] and Amortised Bayesian meta-learning (ABML) [Ravi and Beatson, 2019] and show results for the induction setting on all available testing 5-way 1-shot tasks of mini-ImageNet in Figure 6c. Again, the prediction accuracy made by the model trained on tasks selected by PTM outperforms other baselines, especially the random one by a large margin.

5 CONCLUSION

We propose a generative approach based on variational auto-encoding and LDA adopted in topic modelling to model tasks used in meta-learning. Under this modelling approach, the dataset associated with a task can be expressed as a mixture model of finite Gaussian distributions, where each task differs at the mixture coefficients. An online VI method is presented to infer the parameters of the Gaussian task-theme distributions. The obtained model allows us to represent a task by its variational distribution of mixture coefficient in a latent task-theme simplex, enabling the quantification of either the task uncertainty or task similarity for active task selection.

Acknowledgements

This work was supported by Australian Research Council through grants DP180103232 and FT190100525. We also acknowledge the Phoenix HPC service at the University of Adelaide that provided the super-computing resources for this work.

References

- Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charles C Fowlkes, Stefano Soatto, and Pietro Perona. TASK2VEC: Task embedding for meta-learning. In *International Conference on Computer Vision*, pages 6430–6439, 2019.
- Kelsey Allen, Evan Shelhamer, Hanul Shin, and Joshua Tenenbaum. Infinite mixture prototypes for few-shot learning. In *International Conference on Machine Learning*, pages 232–241. PMLR, 2019.
- Bart Bakker and Tom Heskes. Task clustering and gating for Bayesian multitask learning. *Journal of Machine Learning Research*, 4(May):83–99, 2003.
- Christopher M Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.
- Kevin Canini, Lei Shi, and Thomas Griffiths. Online inference of topics with latent Dirichlet allocation. In *Artificial Intelligence and Statistics*, pages 65–72, 2009.
- Liam Collins, Aryan Mokhtari, and Sanjay Shakkottai. Task-robust model-agnostic meta-learning. In *Advances in Neural Information Processing Systems*, 2020.
- Rajarshi Das, Manzil Zaheer, and Chris Dyer. Gaussian lda for topic models with word embeddings. In *Proceedings of the 53rd Annual Meeting of the Association for*

- Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 795–804, 2015.
- Guneet S Dhillon, Pratik Chaudhari, Avinash Ravichandran, and Stefano Soatto. A baseline for few-shot image classification. In *International Conference on Learning Representations*, 2020.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135, 2017.
- James Foulds, Levi Boyles, Christopher DuBois, Padhraic Smyth, and Max Welling. Stochastic collapsed variational Bayesian inference for latent Dirichlet allocation. In *International Conference on Knowledge Discovery and Data Mining (ACM SIGKDD)*, pages 446–454, 2013.
- Thomas L Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.
- Matthew Hoffman, Francis R Bach, and David M Blei. On-line learning for latent Dirichlet allocation. In *Advances in Neural Information Processing Systems*, pages 856–864, 2010.
- Laurent Jacob, Jean-philippe Vert, and Francis R Bach. Clustered multi-task learning: A convex formulation. In *Advances in Neural Information Processing Systems*, pages 745–752, 2009.
- Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representation*, 2014.
- Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Fei-Fei Li and Pietro Perona. A Bayesian hierarchical model for learning natural scene categories. In *International Conference on Computer Vision and Pattern Recognition*, volume 2, pages 524–531, 2005.
- Gabriel Loaiza-Ganem and John P Cunningham. The continuous Bernoulli: fixing a pervasive error in variational autoencoders. In *Advances in Neural Information Processing Systems*, 2019.
- David JC MacKay. The evidence framework applied to classification networks. *Neural computation*, 4(5):720–736, 1992.
- Thomas Minka. Estimating a Dirichlet distribution, 2000.
- Cuong Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Uncertainty in model-agnostic meta-learning using variational inference. In *Winter Conference on Applications of Computer Vision*, pages 3090–3100, 2020a.
- Cuong V Nguyen, Tal Hassner, Cedric Archambeau, and Matthias Seeger. Leap: A new measure to evaluate transferability of learned representations. In *International Conference on Machine Learning*, 2020b.
- Jonathan K Pritchard, Matthew Stephens, and Peter Donnelly. Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959, 2000.
- Sachin Ravi and Alex Beaton. Amortized bayesian meta-learning. In *International conference on learning representation*, 2019.
- Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International Conference on Learning Representations*, 2017.
- Andrei A Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In *International Conference on Learning Representations*, 2019.
- Paul Ruvolo and Eric Eaton. Active task selection for lifelong machine learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 2013.
- Changjian Shui, Mahdieh Abbasi, Louis-Emile Robitaille, Boyu Wang, and Christian Gagné. A principled approach for learning task similarity in multitask learning. In *International Joint Conference on Artificial Intelligence*, pages 3446–3452, 2019.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.
- Yee W Teh, David Newman, and Max Welling. A collapsed variational Bayesian inference algorithm for latent Dirichlet allocation. In *Advances in neural information processing systems*, pages 1353–1360, 2007.
- Sebastian Thrun and Joseph O’Sullivan. Discovering structure in multiple learning tasks: The TC algorithm. In *International Conference on Machine Learning*, volume 96, pages 489–497, 1996.
- Anh T Tran, Cuong V Nguyen, and Tal Hassner. Transferability and hardness of supervised classification tasks. In *International Conference on Computer Vision*, pages 1395–1405, 2019.
- Amos Tversky. Features of similarity. *Psychological review*, 84(4):327, 1977.

Oriol Vinyals, Charles Blundell, Tim Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, pages 3630–3638, 2016.

Ya Xue, Xuejun Liao, Lawrence Carin, and Balaji Krishnapuram. Multi-task learning for classification with Dirichlet process priors. *Journal of Machine Learning Research*, 8(Jan):35–63, 2007.

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*, pages 7343–7353, 2018.

Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Conference on Computer Vision and Pattern Recognition*, pages 3712–3722, 2018.

Yu Zhang and Dit-Yan Yeung. A convex formulation for learning task relationships in multi-task learning. In *Conference on Uncertainty in Artificial Intelligence*, 2012.

Preliminary version