

# TASK2VEC: Task Embedding for Meta-Learning

Alessandro Achille<sup>1,2</sup> Michael Lam<sup>1</sup> Rahul Tewari<sup>1</sup> Avinash Ravichandran<sup>1</sup>  
 Subhransu Maji<sup>1,3</sup> Charless Fowlkes<sup>1,4</sup> Stefano Soatto<sup>1,2</sup> Pietro Perona<sup>1,5</sup>

achille@cs.ucla.edu {michlam,tewarir,ravinash,smmaji,fowlkec,soattos,peronapp}@amazon.com

<sup>1</sup>AWS <sup>2</sup>University of California, Los Angeles <sup>3</sup>University of Massachusetts, Amherst <sup>4</sup>University of California, Irvine <sup>5</sup>Caltech

## Abstract

*We introduce a method to generate vectorial representations of visual classification tasks which can be used to reason about the nature of those tasks and their relations. Given a dataset with ground-truth labels and a loss function, we process images through a “probe network” and compute an embedding based on estimates of the Fisher information matrix associated with the probe network parameters. This provides a fixed-dimensional embedding of the task that is independent of details such as the number of classes and requires no understanding of the class label semantics. We demonstrate that this embedding is capable of predicting task similarities that match our intuition about semantic and taxonomic relations between different visual tasks. We demonstrate the practical value of this framework for the meta-task of selecting a pre-trained feature extractor for a novel task. We present a simple meta-learning framework for learning a metric on embeddings that is capable of predicting which feature extractors will perform well on which task without actually fine-tuning the model. Selecting a feature extractor with task embedding yields performance close to the best available feature extractor, with substantially less computational effort than exhaustively training and evaluating all available models.*

## 1. Introduction

The success of deep learning in computer vision is due in part to the fact that models trained for one task can often be used on related tasks. Yet, no general framework exists to describe and reason about relations between tasks. We introduce the TASK2VEC embedding, a technique to represent tasks as elements of a vector space based on the Fisher Information Matrix. The norm of the embedding correlates with the complexity of the task, while the distance between embeddings captures semantic similarities between tasks (Fig. 1). When other natural distances are available, such as the taxonomic distance in biological classification, we find they correlate well with the embedding distance (Fig. 2). We also introduce an asymmetric distance on the embed-

ding space that correlates with transferability between tasks.

Computation of the embedding leverages a duality between parameters (weights) and outputs (activations) in a deep neural network (DNN). Just as the activations of a DNN trained on a complex visual recognition task are a rich representation of the input images, we show that the gradients of the weights relative to a task-specific loss are a rich representation of the task itself. Given a task defined by a dataset  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$  of labeled samples, we feed the data through a pre-trained reference convolutional neural network which we call a “probe network”, and compute the diagonal Fisher Information Matrix (FIM) of the network filter parameters to capture the structure of the task (Sect. 3). Since the architecture and weights of the probe network are fixed, the FIM provides a fixed-dimensional representation of the task which is independent of, e.g., how many categories there are. We show that this embedding simultaneously encodes the “difficulty” of the task, statistics of the input domain, and which features extracted by the probe network are discriminative for the task (Sect. 3.2).

Our task embedding can be used to reason about the space of tasks and solve meta-tasks. As a motivating example, we study the problem of selecting the best pre-trained feature extractor to solve a new task (Sect. 4). This is particularly valuable when there is insufficient data to train or fine-tune a generic model, and transfer of knowledge is essential. To select an appropriate pre-trained model, we design a joint embedding of models and tasks in the same vector space, which we call MODEL2VEC. We formulate this as a meta-learning problem where the objective is to find an embedding such that that models whose embeddings are close to a task exhibit good performance on that task.

We present large-scale experiments on a library of 1,460 fine-grained classification tasks constructed from existing computer vision datasets. These tasks vary in the level of difficulty and have orders of magnitude variation in training set size, mimicking the heavy-tailed distribution of real-world tasks. Our experiments show that using TASK2VEC to select an expert from a collection of 156 feature extractors outperforms the standard practice of fine-tuning a generic

model trained on ImageNet. We find that the selected expert is close to optimal while being orders of magnitude faster than brute-force selection.

## 2. Background and Related Work

What metric should be used on the space of tasks? This depends critically on the meta-task we are considering. For the purpose of model selection there are several natural metrics that may be considered.

**Domain distance.** Tasks distinguished by their domain can be understood simply in terms of image statistics. Due to the bias of different datasets, sometimes a benchmark task can be identified just by looking at a few images [34]. The question of determining what summary statistics are useful (analogous to our choice of probe network) has also been considered. For example [9] train an autoencoder that learns to extract fixed dimensional summary statistics that can reproduce many different datasets accurately. However, for general vision tasks, the domain is insufficient (detecting pedestrians and reading license plates are different tasks that share the same domain of street scene images).

**Taxonomic distance.** Some collections of tasks come with a natural notion of semantic similarity based on a taxonomic hierarchy. We may say that classifying dog breeds is closer to classification of cats than it is to classification of plant species. When each task is specified by a set of categories in a subtree of the hierarchy, the tasks inherit a distance from the taxonomy. In this setting, we can define

$$D_{\text{tax}}(t_a, t_b) = \min_{i \in S_a, j \in S_b} d(i, j),$$

where  $S_a, S_b$  are the sets of categories in task  $t_a, t_b$  and  $d(i, j)$  is an ultrametric or graph distance on the taxonomy. However such taxonomies are not available for all tasks and visual similarity need not be correlated with semantic similarity.

**Transfer distance.** Another notion that abstracts away the details of domains and labels is the transfer distance between task. For a specified DNN architecture, this is defined as the gain in performance from pre-training a model on task  $t_a$  and then fine-tuning for task  $t_b$  relative to the performance of simply training on task  $t_b$  using a fixed initialization (random or generic pre-trained)<sup>1</sup>. We write

$$D_{\text{ft}}(t_a \rightarrow t_b) = \frac{\mathbb{E}[\ell_{a \rightarrow b}] - \mathbb{E}[\ell_b]}{\mathbb{E}[\ell_b]},$$

where the expectations are taken over all training runs with the selected architecture, training procedure and network initialization,  $\ell_b$  is the final test error obtained by training

<sup>1</sup>We improperly call this a distance but it is not necessarily symmetric or even positive.

on task  $b$  from the chosen initialization, and  $\ell_{a \rightarrow b}$  is the error obtained instead when starting from a solution to task  $a$  and then fine-tuning (with the selected procedure) on task  $t_b$ .

This notion of task transfer is the focus of Taskonomy [39], which explores knowledge transfer in a curated collection of 26 visual tasks, ranging from classification to 3D reconstruction, defined on a common domain. They compute transfer distances between pairs of tasks and use the results to compute a directed hierarchy. Adding novel tasks to the hierarchy requires computing the transfer distance to all other tasks in the collection.

In contrast, we show it is possible to directly produce a task embedding in constant time without computing pairwise distances. This makes it feasible to experiment on a much larger library of 1,460 classification tasks in multiple domains. The large task collection and cheap cost of embedding allows us to tackle new meta-learning problems.

**Fisher Kernels and Fisher Information.** Our work takes inspiration from Jaakkola and Hausler [16]. They propose the ‘‘Fisher Kernel’’, which uses the gradients of a generative model score function as a representation of similarity between data items

$$K(x^{(1)}, x^{(2)}) = \nabla_{\theta} \log P(x^{(1)}|\theta)^T F^{-1} \nabla_{\theta} \log P(x^{(2)}|\theta).$$

Here,  $P(x|\theta)$  is a parameterized generative model and  $F$  is the Fisher information matrix. This provides a way to utilize generative models in the context of discriminative learning. Variants of the Fisher kernel have found wide use as a representation of images [28, 29] and other structured data such as protein molecules [17] and text [30]. Since the generative model can be learned on unlabeled data, several works have investigated the use of Fisher kernel for unsupervised learning [14, 31]. [35] learns a metric on the Fisher kernel representation similar to our metric learning approach. Our approach differs in that we use the FIM as a representation of a whole dataset (task) rather than using model gradients as representations of individual data items.

**Fisher Information for CNNs.** Our approach to task embedding makes use of the Fisher Information matrix of a neural network as a characterization of the task. Use of Fisher information for neural networks was popularized by Amari [6] who advocated optimization using natural gradient descent which leverages the fact that the FIM is a parameterization-independent metric on statistical models. Recent work has focused on approximators of FIM appropriate in this setting (see e.g., [12, 10, 25]). FIM has also been proposed for various regularization schemes [5, 8, 22, 27], to analyze learning dynamics of deep networks [4], and to overcome catastrophic forgetting [19].

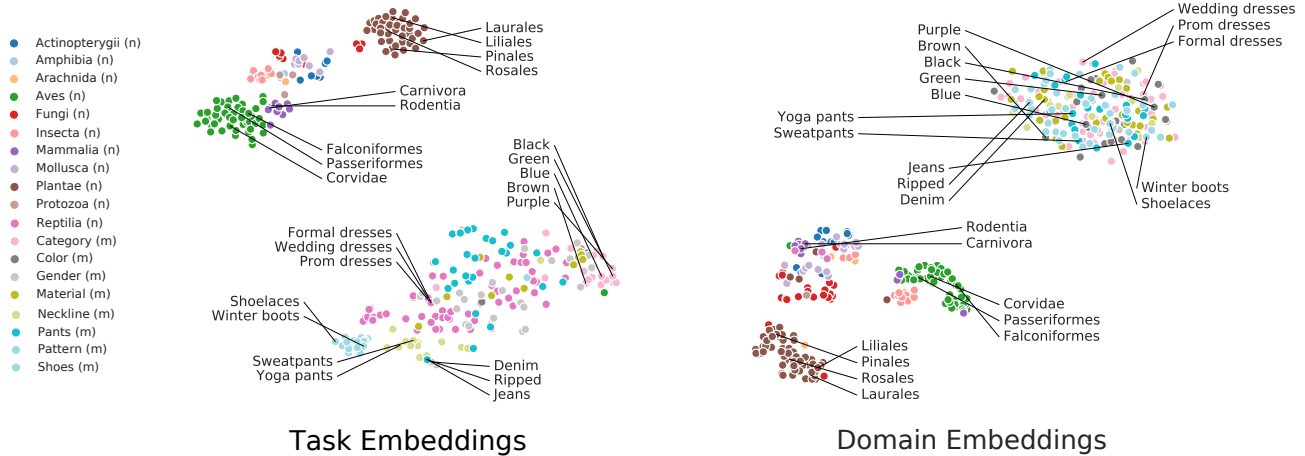


Figure 1: **Task embedding across a large library of tasks** (best seen magnified). **(Left)** T-SNE visualization of the embedding of tasks extracted from the iNaturalist, CUB-200, iMaterialist datasets. Colors indicate ground-truth grouping of tasks based on taxonomic or semantic types. Notice that the bird classification tasks extracted from CUB-200 embed near the bird classification task from iNaturalist, even though the original datasets are different. iMaterialist is well separated from iNaturalist, as it entails very different tasks (clothing attributes). Notice that some tasks of similar type (such as color attributes) cluster together but attributes of different task types may also mix when the underlying visual semantics are correlated. For example, the tasks of jeans (clothing type), denim (material) and ripped (style) recognition are close in the task embedding. **(Right)** T-SNE visualization of the domain embeddings (using mean feature activations) for the same tasks. Domain embedding can distinguish iNaturalist tasks from iMaterialist tasks due to differences in the two problem domains. However, the fashion attribute tasks on iMaterialist all share the same domain and only differ in their labels. In this case, the domain embeddings collapse to a region without recovering any sensible structure.

**Meta Learning and Model Selection.** Meta-learning has a long history with much recent work dedicated to meta-tasks such as neural architecture search, hyper-parameter estimation and robust few-shot learning. The meta-learning problem of model selection seeks to choose from a library of classifiers to solve a new task [33, 2, 20]. Unlike our approach, these previous techniques usually address the question via land-marking or active testing, in which a few different models are evaluated and performance of the remainder estimated by extension. This can be viewed as a problem of completing unknown entries in a matrix defined by performance of each model on each task. In computer vision, this idea has been explored for selecting a detector for a new category out of a large library of detectors [26, 40, 38].

### 3. Task Embeddings via Fisher Information

Given an observed input image  $x$  and an unknown task variable  $y$  (e.g., a label), a deep network is a family of functions  $p_w(y|x)$  parametrized by weights  $w$ , trained to approximate the posterior  $p(y|x)$  by minimizing the (possibly regularized) cross entropy loss  $H_{p_w, \hat{p}}(y|x) = \mathbb{E}_{x, y \sim \hat{p}}[-\log p_w(y|x)]$ , where  $\hat{p}$  is the empirical distribution defined by the training set  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ . It is useful, especially in transfer learning, to think of the net-

work as composed of two parts: a feature extractor which computes some representation  $z = \phi_w(x)$  of the input data, and a “head,” or classifier, which predicts the distribution  $p(y|z)$  given the representation  $z$ .

Not all network weights are equally important in predicting the task variable. The importance, or “informative content”, of a weight for the task can be quantified by considering a perturbation  $w' = w + \delta w$  of the weights, and measuring the average Kullback-Leibler (KL) divergence between the original output distribution  $p_w(y|x)$  and the perturbed one  $p_{w'}(y|x)$ . To second-order approximation, this is

$$\mathbb{E}_{x \sim \hat{p}} KL(p_{w'}(y|x) \| p_w(y|x)) = \delta w^T \cdot F \cdot \delta w + o(\delta w^2),$$

where  $F$  is the Fisher information matrix (FIM):

$$F = \mathbb{E}_{x, y \sim \hat{p}(x)p_w(y|x)} [\nabla_w \log p_w(y|x) \nabla_w \log p_w(y|x)^T].$$

that is, the covariance of the scores (gradients of the log-likelihood) with respect to the model parameters.

The FIM is a Riemannian metric on the space of probability distributions [7], and provides a measure of the information a particular parameter (weight or feature) contains about the joint distribution  $p_w(x, y) = p_w(y|x)\hat{p}(x)$ . If the classification performance for a given task does not depend strongly on a parameter, the corresponding entries in the

FIM will be small. The FIM is also related to the (Kolmogorov) complexity of a task, a property that can be used to define a computable metric of the learning distance between tasks [3]. Finally, the FIM can be interpreted as an easy-to-compute positive semidefinite upper-bound to the Hessian of the cross-entropy loss and coincides with it at local minima [24]. In particular, “flat minima” correspond to weights that have, on average, low Fisher information [5, 13].

### 3.1. TASK2VEC embedding using a probe network

While the network activations capture the information in the input image which are needed to infer the image label, the FIM indicates the set of feature maps which are more informative for solving the current task. Following this intuition, we use the FIM to represent the task itself. However, the FIMs computed on different networks are not directly comparable. To address this, we use a single “probe” network pre-trained on ImageNet as a feature extractor and re-train only the classifier layer on any given task, which usually can be done efficiently. After training is complete, we compute the FIM for the feature extractor parameters.

Since the full FIM is unmanageably large for rich probe networks based on CNNs, we make two additional approximations. First, we only consider the diagonal entries, which implicitly assumes that correlations between different filters in the probe network are not important. Second, since the weights in each filter are usually not independent, we average the Fisher Information for all weights in the same filter. The resulting representation thus has fixed size, equal to the number of filters in the probe network. We call this embedding method TASK2VEC.

**Robust Fisher computation.** Since the FIM is a local quantity, it is affected by the local geometry of the training loss landscape, which is highly irregular in many deep network architectures [21], and may be too noisy when trained with few samples. To avoid this problem, instead of direct computation, we use a more robust estimator that leverages connections to variational inference. Assume we perturb the weights  $\hat{w}$  of the network with Gaussian noise  $\mathcal{N}(0, \Lambda)$  with precision matrix  $\Lambda$ , and we want to find the optimal  $\Lambda$  which yields a good expected error, while remaining close to an isotropic prior  $\mathcal{N}(\hat{w}, \lambda^2 I)$ . That is, we want to find  $\Lambda$  that minimizes:

$$L(\hat{w}; \Lambda) = \mathbb{E}_{w \sim \mathcal{N}(\hat{w}, \Lambda)} [H_{p_w, \hat{p}}(y|x)] + \beta KL(\mathcal{N}(0, \Lambda) \parallel \mathcal{N}(0, \lambda^2 I)),$$

where  $H$  is the cross-entropy loss and  $\beta$  controls the weight of the prior. Notice that for  $\beta = 1$  this reduces to the Evidence Lower-Bound (ELBO) commonly used in variational inference. Approximating to the second order, the optimal

value of  $\Lambda$  satisfies (see Supplementary Material):

$$\frac{\beta}{2N} \Lambda = F + \frac{\beta \lambda^2}{2N} I.$$

Therefore,  $\frac{\beta}{2N} \Lambda \sim F + o(1)$  can be considered as an estimator of the FIM  $F$ , biased towards the prior  $\lambda^2 I$  in the low-data regime instead of being degenerate. In case the task is trivial (the loss is constant or there are too few samples) the embedding will coincide with the prior  $\lambda^2 I$ , which we will refer to as the *trivial embedding*. This estimator has the advantage of being easy to compute by directly minimizing the loss  $L(\hat{w}; \Sigma)$  through Stochastic Gradient Variational Bayes [18], while being less sensitive to irregularities of the loss landscape than direct computation, since the value of the loss depends on the cross-entropy in a neighborhood of  $\hat{w}$  of size  $\Lambda^{-1}$ . As mentioned previously, we estimate one parameter per filter, rather than per weight, which in practice means that we constrain  $\Lambda_{ii} = \Lambda_{jj}$  whenever  $w_i$  and  $w_j$  belongs to the same filter. In this case, optimization of  $L(\hat{w}; \Lambda)$  can be done efficiently using the local reparametrization trick of [18].

### 3.2. Properties of the TASK2VEC embedding

The task embedding we just defined has a number of useful properties. For illustrative purposes, consider a two-layer sigmoidal network for which an analytic expression can be derived (see Supplementary Materials). The FIM of the feature extractor parameters can be written using the Kronecker product as

$$F = \mathbb{E}_{x, y \sim \hat{p}(x)p_w(y|x)} [(y - p)^2 \cdot S \otimes xx^T]$$

where  $p = p_w(y = 1|x)$  and the matrix  $S = ww^T \odot zz^T \odot (1 - z)(1 - z)^T$  is an element-wise product of classifier weights  $w$  and first layer feature activations  $z$ . It is informative to compare this expression to an embedding based only on the dataset domain statistics, such as the (non-centered) covariance  $C_0 = \mathbb{E}[xx^T]$  of the input data or the covariance  $C_1 = \mathbb{E}[zz^T]$  of the feature activations. One could take such statistics as a representative *domain embedding* since they only depend on the marginal distribution  $p(x)$  in contrast to the FIM *task embedding*, which depends on the joint distribution  $p(x, y)$ . These simple expressions highlight some important (and more general) properties of the Fisher embedding we now describe.

**Invariance to the label space.** The task embedding does not directly depend on the task labels but only on the predicted distribution  $p_w(y|x)$  of the trained model. Information about the ground-truth labels  $y$  is encoded in the weights  $w$  which are a sufficient statistic of the task [5]. In particular, the task embedding is invariant to permutations of the labels  $y$  and has fixed dimension (number of filters of the feature extractor) regardless of the output space (e.g., k-way classification with varying k).

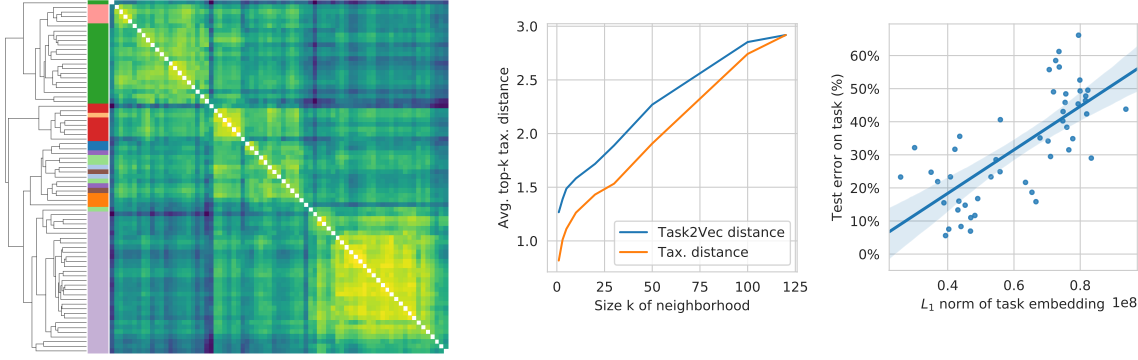


Figure 2: **Distance between species classification tasks.** (Left) Task similarity matrix ordered by hierarchical clustering. Note that the dendrogram produced by the task similarity matches the taxonomic clusters (indicated by color bar). (Center) For tasks extracted from iNaturalist and CUB, we compare the cosine distance between tasks to their taxonomical distance. As the size of the task embedding neighborhood increases (measured by number of tasks in the neighborhood), we plot the average taxonomical distance of tasks from the neighborhood center. While the task distance does not perfectly match the taxonomical distance (whose curve is shown in orange), it shows a good correlation. Difference are both due to the fact that taxonomically close species may need very different features to be classified, creating a mismatch between the two notions of distance, and because for some tasks in iNaturalist too few samples are provided to compute a good embedding. (Right) Correlation between  $L_1$  norm of the task embedding (distance from origin) and test error obtained on the task.

**Encoding task difficulty.** As we can see from the expressions above, if the model is very confident in its predictions,  $\mathbb{E}[(y - p)^2]$  goes to zero. Hence, the norm of the task embedding  $\|F\|_*$  scales with the difficulty of the task for a given feature extractor  $\phi$ . Fig. 2 (Right) shows that even for more complex models trained on real data, the FIM norm correlates with test performance.

**Encoding task domain.** Data points  $x$  that are classified with high confidence, i.e.,  $p$  is close to 0 or 1, will have a lower contribution to the task embedding than points near the decision boundary since  $p(1 - p)$  is maximized at  $p = 1/2$ . Compare this to the covariance matrix of the data,  $C_0$ , to which all data points contribute equally. Instead, in TASK2VEC the information on the domain is based on data near the decision boundary (task-weighted domain embedding).

**Encoding useful features for the task.** The FIM depends on the curvature of the loss function with the diagonal entries capturing the sensitivity of the loss to model parameters. Specifically, in the two-layer model, one can see that if a given feature is uncorrelated with  $y$ , the corresponding blocks of  $F$  are zero. In contrast, a domain embedding based on feature activations of the probe network (e.g.,  $C_1$ ) only reflects which features vary over the dataset without indication of whether they are relevant to the task.

### 3.3. Symmetric and asymmetric TASK2VEC metrics

By construction, the Fisher embedding on which TASK2VEC is based captures fundamental information

about the structure of the task. We may therefore expect that the distance between two embeddings correlates positively with natural metrics on the space of tasks. However, there are two problems in using the Euclidean distance between embeddings: the parameters of the network have different scales, and the norm of the embedding is affected by complexity of the task and the number of samples used to compute the embedding.

**Symmetric TASK2VEC distance.** To make the distance computation robust, we propose to use the cosine distance between normalized embeddings:

$$d_{\text{sym}}(F_a, F_b) = d_{\text{cos}}\left(\frac{F_a}{F_a + F_b}, \frac{F_b}{F_a + F_b}\right),$$

where  $d_{\text{cos}}$  is the cosine distance,  $F_a$  and  $F_b$  are the two task embeddings (i.e., the diagonal of the Fisher Information computed on the same probe network), and the division is element-wise. This is a symmetric distance which we expect to capture semantic similarity between two tasks. For example, we show in Fig. 2 that it correlates well with the taxonomical distance between species on iNaturalist.

On the other hand, precisely for this reason, this distance is ill-suited for tasks such as model selection, where the (inherently asymmetric) transfer distance is more relevant.

**Asymmetric TASK2VEC distance.** In a first approximation, that does not consider either the model or the training procedure used, positive transfer between two tasks depends both on the similarity between two tasks and on the complexity of the first. Indeed, pre-training on a general but

complex task such as ImageNet often yields a better result than fine-tuning from a close dataset of lower complexity. In our case, complexity can be measured as the distance from the trivial embedding. This suggests the following asymmetric score, again improperly called a “distance” despite being asymmetric and possibly negative:

$$d_{\text{asym}}(t_a \rightarrow t_b) = d_{\text{sym}}(t_a, t_b) - \alpha d_{\text{sym}}(t_a, t_0),$$

where  $t_0$  is the trivial embedding, and  $\alpha$  is an hyperparameter. This has the effect of bringing more complex models closer. The hyper-parameter  $\alpha$  can be selected based on the meta-task. In our experiments, we found that the best value of  $\alpha$  ( $\alpha = 0.15$  when using a ResNet-34 pretrained on ImageNet as the probe network) is robust to the choice of meta-tasks.

### 3.4. MODEL2VEC: Co-embedding models and tasks

By construction, the TASK2VEC distance ignores details of the model and only relies on the task. If we know what task a model was trained on, we can represent the model by the embedding of that task. However, in general we may not have such information (e.g., black-box models or hand-constructed feature extractors). We may also have multiple models trained on the same task with different performance characteristics. To model the joint interaction between task and model (i.e., architecture and training algorithm), we aim to learn a joint embedding of the two.

We consider for concreteness the problem of learning a joint embedding for model selection. In order to embed models in the task space so that those near a task are likely to perform well on that task, we formulate the following meta-learning problem: Given  $k$  models, their MODEL2VEC embedding are the vectors  $m_i = F_i + b_i$ , where  $F_i$  is the task embedding of the task used to train model  $m_i$  (if available, else we set it to zero), and  $b_i$  is a learned “model bias” that perturbs the task embedding to account for particularities of the model. We learn  $b_i$  by optimizing a  $k$ -way cross entropy loss to predict the best model given the task distance (see Supplementary Material):

$$\mathcal{L} = \mathbb{E}[-\log p(m \mid d_{\text{asym}}(t, m_0), \dots, d_{\text{asym}}(t, m_k))].$$

After training, given a novel query task  $t$ , we can then predict the best model for it as the  $\arg \min_i d_{\text{asym}}(t, m_i)$ , that is, the model  $m_i$  embedded closest to the query task.

## 4. Experiments

We test TASK2VEC on a large collection of tasks and models with a range of task similarities<sup>2</sup>. Our experiments aim to test both qualitative properties of the embedding and

<sup>2</sup>Here we discuss only classification tasks, the supplement describes additional experiments on pixel-labeling and regression tasks

its performance on meta-learning tasks. We use an off-the-shelf ResNet-34 pretrained on ImageNet as our probe network, which we found to give the best overall performance (see Sect. 4.2). The collection of tasks is generated starting from the following four main datasets. **iNaturalist** [36]: Each task corresponds to species classification in a given taxonomical order. For instance, the “*Rodentia* task” is to classify species of rodents. Each task is defined on a separate subset of the images in the original dataset; that is, the domains of the tasks are disjoint. **CUB-200** [37]: We use the same procedure as iNaturalist to create tasks. All tasks are classifications inside orders of birds (the *aves* taxonomical class) and have generally much fewer training samples than corresponding tasks in iNaturalist. **iMaterialist** [1] and **DeepFashion** [23]: Each image in these datasets is associated with several binary attributes (e.g., style attributes) and categorical attributes (e.g., color, type of dress, material). We binarize the categorical attributes and consider each attribute as a separate task. Notice that in this case, all tasks share the same domain and are naturally correlated.

In total, our collection of tasks has 1460 tasks (207 iNaturalist, 25 CUB, 228 iMaterialist, 1000 DeepFashion). While a few tasks have many training examples (e.g., hundred thousands), most have just hundreds or thousands of samples. This simulates the heavy-tail distribution of data in real-world applications.

For model selection experiments, we assemble a library of “expert” feature extractors. These are ResNet-34 models pre-trained on ImageNet and then fine-tuned on a specific task or collection of related tasks (see Supplementary Materials for details). We also consider a “generic” expert pre-trained on ImageNet without any fine-tuning. Finally, for each combination of expert feature extractor and task, we trained a linear classifier on top of the expert in order to solve the selected task using the expert.

In total, we trained 4,100 classifiers, 156 feature extractors and 1,460 embeddings. The total effort to generate the final results was about 1,300 GPU hours.

**Meta-tasks.** In Sect. 4.2, for a given task we aim to predict, using TASK2VEC, which expert feature extractor will yield the best classification performance. In particular, we formulate two model selection meta-tasks: **iNat + CUB** and **Mixed**. The first consists of 50 tasks and experts from iNaturalist and CUB, and aims to test fine-grained expert selection in a restricted domain. The second contains a mix of 26 curated experts and 50 random tasks extracted from all datasets, and aims to test model selection between different domains and tasks (see Supplementary Material for details).

### 4.1. Task Embedding Results

**Task Embedding qualitatively reflects taxonomic distance for iNaturalist.** For tasks extracted from the iNaturalist dataset (classification of species), the taxonomical



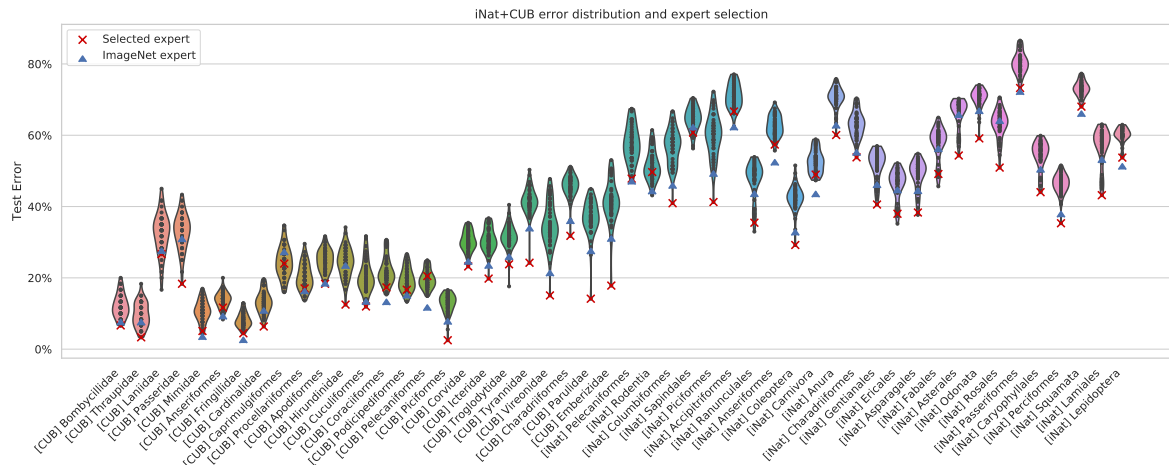


Figure 3: **TASK2VEC often selects the best available experts.** Violin plot of the test error distribution (shaded) on tasks from the CUB-200 dataset (columns) obtained by training a linear classifier over several expert feature extractors (points). Most specialized feature extractors perform similarly on a given task, and similar or worse than a generic feature extractor pre-trained on ImageNet (blue triangles). However, in some cases a carefully chosen expert, trained on a related task, can greatly outperform all others (long lower whiskers). The model selection algorithm based on TASK2VEC can predict an expert to use for the task (red cross, lower is better) and often recommends the optimal, or near optimal, feature extractor without performing an expensive brute-force training and evaluation over all available experts. Columns are ordered by norm of the task embedding vector. Tasks with lower embedding norm have lower error and more “complex” task (task with higher embedding norm) tend to benefit more from a specialized expert.

distance between orders provides a natural metric of the semantic similarity between tasks. In Fig. 2, we compare the symmetric TASK2VEC distance with the taxonomical distance, showing strong agreement.

**Task embedding for iMaterialist.** In Fig. 1, we show a t-SNE visualization of the embedding for iMaterialist and iNaturalist tasks. Task embedding yields interpretable results: Tasks that are correlated in the dataset, such as binary classes corresponding to the same categorical attribute, may end up far away from each other and close to other tasks that are semantically more similar (e.g., the *jeans* category task is close to the *ripped* attribute and the *denim* material). In the visualization, this non-trivial grouping is reflected in the mixture of colors of semantically related nearby tasks.

We also compare the TASK2VEC embedding with a domain embedding baseline, which only exploits the input distribution  $p(x)$  rather than the task distribution  $p(x, y)$ . While some tasks are highly correlated with their domain (e.g., tasks from iNaturalist), other tasks differ only on the labels (e.g., all the attribute tasks of iMaterialist, which share the same clothes domain). Accordingly, the domain embedding recovers similar clusters on iNaturalist. However, on iMaterialist, domain embedding collapses all tasks to a single uninformative cluster (not a single point due to slight noise in embedding computation).

**Task Embedding encodes task difficulty.** The scatter-plot in Fig. 1 compares the norm of embedding vectors vs. per-

formance of the best expert (or task specific model for cases where pre-train and test tasks coincide). As suggested by the analysis for the two-layer model, the norm of the task embedding also correlates with the complexity of the task for real tasks and model architectures.

## 4.2. Model Selection

Given a task, our aim is to select an expert feature extractor that maximizes the classification performance on that task. We propose two strategies: (1) embed the task and select the feature extractor trained on the most similar task, and (2) jointly embed the models and tasks, and select a model using the learned metric (see Section 3.4). Notice that (1) does not use knowledge of the model performance on various tasks, which makes it more widely applicable but requires we know what task a model was trained for and may ignore the fact that models trained on slightly different tasks may still provide an overall better feature extractor (for example by over-fitting less to the task they were trained on).

In Table 1 we compare the overall results of the various proposed metrics on the model selection meta-tasks. On both the iNat+CUB and Mixed meta-tasks, the Asymmetric TASK2VEC model selection is close to the ground-truth optimal, and significantly improves over both chance, and over using a generic ImageNet expert. Notice that our method has  $O(1)$  complexity, while searching over a collection of

Meta-task	Optimal	Chance	ImageNet	TASK2VEC	Asymmetric TASK2VEC	MODEL2VEC
iNat + CUB	31.24	+59.52%	+30.18%	+42.54%	+9.97%	<b>+6.81%</b>
Mixed	22.90	+112.49%	+75.73%	+40.30%	+29.23%	<b>+27.81%</b>

Table 1: **Model selection performance.** Average error obtained on two meta-learning datasets by exhaustive search to select the optimal expert, and relative error increase when using cheaper model selection methods. A general model pre-trained on ImageNet performs better than picking an expert at random (chance). However, picking an expert using the Asymmetric TASK2VEC distance achieves substantially better performance and can be further improved by meta-learning (MODEL2VEC).

$N$  experts is  $O(N)$ .

**Error distribution.** In Fig. 3, we show in detail the error distribution of the experts on multiple tasks. It is interesting to observe that the classification error obtained using most experts clusters around some mean value, and little improvement is observed over using a generic expert. On the other hand, for each task there often exist a few experts that can obtain significantly better performance on the task than a generic expert. This confirms the importance of having access to a large collection of experts when solving a new task, especially if few training data are available. TASK2VEC provides an efficient way to index this collection and identify an appropriate expert for a new task without brute-force search.

**Dependence on task dataset size.** Selecting pre-trained experts is especially important when the novel query task has relatively few training samples. In Fig. 4, we show how the performance of TASK2VEC model selection varies as a function of dataset size. Even for tasks with few samples, TASK2VEC selection performs nearly as well as using the optimal (ground-truth) expert. If in addition to training a classifier, we fine-tune the selected expert, error decreases further. At all dataset sizes, we see that the expert selected by TASK2VEC significantly outperforms the standard baseline of using a generic ImageNet pre-trained model.

**Choice of probe network.** In Table 2, we show that DenseNet [15] and ResNet architectures [11] perform significantly better when used as probe networks to compute the TASK2VEC embedding than a VGG [32] architecture.

## 5. Discussion

TASK2VEC is an efficient way to represent a task as a fixed dimensional vector with several appealing properties.

	Chance	VGG-13	DenseNet-121	ResNet-13
Top-10	+13.95%	+4.82%	+0.30%	<b>+0.00%</b>
All	+59.52%	+38.03%	+10.63%	<b>+9.97%</b>

Table 2: **Choice of probe network.** Mean increase in error relative to optimal (ground-truth) expert selection on the iNat+CUB meta-task for different choices of the probe-network. We also report the performance on a subset of 10 tasks with the most training samples to show the effect of data size in choosing a probe architecture.

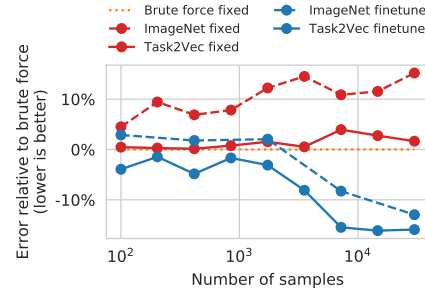


Figure 4: **TASK2VEC improves outperforms baselines at different dataset sizes:** Performance of model selection on a subset of 4 tasks as a function of the number of samples available to train relative to optimal model selection (dashed orange). Training a classifier on the fixed feature extractor selected by TASK2VEC (solid red) is always better than using a generic ImageNet feature extractor (dashed red). The same holds when fine-tuning the feature extractor (blue curves). In the low-data regime, the fixed pre-trained experts selected by TASK2VEC even outperform costly fine-tuning of a generic ImageNet feature extractor (dashed blue).

The embedding norm correlates with the task difficulty, and the cosine distance between embeddings is predictive of both natural semantic distances between tasks (e.g., the taxonomic distance for species classification) and the fine-tuning distance for transfer learning. Having a representation of tasks paves the way for a wide variety of meta-learning tasks. In this work, we focused on selection of an expert feature extractor in order to solve a new task, and showed that using TASK2VEC to select an expert from a collection can improve test performance over the de facto baseline of using an ImageNet pre-trained model while adding only a small overhead to the training process.

We demonstrate that TASK2VEC is scalable to thousands of tasks, allowing us to reconstruct a topology on the task space, and to test meta-learning solutions. The current experiments highlight the usefulness of our methods. Even so, our collection does not capture the full complexity and variety of tasks that one may encounter in real-world situations. Future work should further test effectiveness, robustness, and limitations of the embedding on larger and more diverse collections.



## References

- [1] iMaterialist Fashion Challenge at FGVC5 workshop, CVPR. <https://www.kaggle.com/c/imaterialist-challenge-fashion-2018>. 6
- [2] Salisu Mamman Abdulrahman, Pavel Brazdil, Jan N van Rijn, and Joaquin Vanschoren. Speeding up algorithm selection using average ranking and active testing by introducing runtime. *Machine learning*, 107(1):79–108, 2018. 3
- [3] Alessandro Achille, Glen Mbeng, Giovanni Paolini, and Stefano Soatto. The dynamic distance between learning tasks: From Kolmogorov complexity to transfer learning via quantum physics and the information bottleneck of the weights of deep networks. *Proc. of the NIPS Workshop on Integration of Deep Learning Theories (ArXiv: 1810.02440)*, October 2018. 4
- [4] Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep neural networks. *Proc. of the Intl. Conf. on Learning Representations (ICLR)*. ArXiv:1711.08856, 2019. 2
- [5] Alessandro Achille and Stefano Soatto. Emergence of invariance and disentanglement in deep representations. *Journal of Machine Learning Research (ArXiv 1706.01350)*, 19(50):1–34, 2018. 2, 4
- [6] Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998. 2
- [7] Shun-Ichi Amari and Hiroshi Nagaoka. Methods of information geometry, volume 191 of translations of mathematical monographs. *American Mathematical Society*, 13, 2000. 3
- [8] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. *arXiv preprint arXiv:1802.05296*, 2018. 2
- [9] Harrison Edwards and Amos Storkey. Towards a neural statistician. *arXiv preprint arXiv:1606.02185*, 2016. 2
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 8
- [12] Tom Heskes. On natural learning and pruning in multi-layered perceptrons. *Neural Computation*, 12(4):881–901, 2000. 2
- [13] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997. 4
- [14] Alex D Holub, Max Welling, and Pietro Perona. Combining generative models and fisher kernels for object recognition. In *IEEE International Conference on Computer Vision*, volume 1, pages 136–143. IEEE, 2005. 2
- [15] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 8
- [16] Tommi Jaakkola and David Haussler. Exploiting generative models in discriminative classifiers. In *Advances in Neural Information Processing Systems*, pages 487–493, 1999. 2
- [17] Tommi S Jaakkola, Mark Diekhans, and David Haussler. Using the Fisher kernel method to detect remote protein homologies. In *ISMB*, volume 99, pages 149–158, 1999. 2
- [18] Diederik P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583, 2015. 4
- [19] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, page 201611835, 2017. 2
- [20] Rui Leite, Pavel Brazdil, and Joaquin Vanschoren. Selecting classification algorithms with active testing. In *International workshop on machine learning and data mining in pattern recognition*, pages 117–131. Springer, 2012. 3
- [21] Hao Li, Zheng Xu, Gavin Taylor, and Tom Goldstein. Visualizing the loss landscape of neural nets. *arXiv preprint arXiv:1712.09913*, 2017. 4
- [22] Tengyuan Liang, Tomaso Poggio, Alexander Rakhlin, and James Stokes. Fisher-Rao metric, geometry, and complexity of neural networks. *arXiv preprint arXiv:1711.01530*, 2017. 2
- [23] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1096–1104, 2016. 6
- [24] James Martens. New perspectives on the natural gradient method. *CoRR*, abs/1412.1193, 2014. 4
- [25] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning (ICML)*, pages 2408–2417, 2015. 2
- [26] Pyry Matikainen, Rahul Sukthankar, and Martial Hebert. Model recommendation for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2256–2263. IEEE, 2012. 3
- [27] Youssef Mroueh and Tom Sercu. Fisher GAN. In *Advances in Neural Information Processing Systems*, pages 2513–2523, 2017. 2
- [28] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. Improving the fisher kernel for large-scale image classification. In *European Conference on Computer Vision*, pages 143–156. Springer, 2010. 2
- [29] Jorge Sánchez, Florent Perronnin, Thomas Mensink, and Jakob Verbeek. Image classification with the Fisher vector: Theory and practice. *International Journal of Computer Vision*, 105(3):222–245, 2013. 2
- [30] Craig Saunders, Alexei Vinokourov, and John S Shawe-taylor. String kernels, Fisher kernels and finite state automata. In *Advances in Neural Information Processing Systems*, pages 649–656, 2003. 2
- [31] Matthias Seeger. Learning with labeled and unlabeled data. Technical Report EPFL-REPORT-161327, Institute for Adaptive and Neural Computation, University of Edinburgh, 2000. 2

- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 8
- [33] Michael R Smith, Logan Mitchell, Christophe Giraud-Carrier, and Tony Martinez. Recommending learning algorithms and their associated hyperparameters. *arXiv preprint arXiv:1407.1890*, 2014. 3
- [34] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1521–1528. IEEE, 2011. 2
- [35] Laurens Van Der Maaten. Learning Discriminative Fisher Kernels. In *International Conference on Machine Learning*, volume 11, pages 217–224, 2011. 2
- [36] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018. 6
- [37] 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. 6
- [38] Yu-Xiong Wang and Martial Hebert. Model recommendation: Generating object detectors from few samples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1619–1628, 2015. 3
- [39] Amir R Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3712–3722, 2018. 2
- [40] Peng Zhang, Jiuling Wang, Ali Farhadi, Martial Hebert, and Devi Parikh. Predicting failures of vision systems. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3566–3573, 2014. 3