

The information complexity of learning tasks, their structure and their distance

ALESSANDRO ACHILLE[†]

Department of Computer Science, University of California, Los Angeles, CA, USA

[†]Corresponding author. Email: achille@cs.ucla.edu

GIOVANNI PAOLINI

Classe di Scienze, Scuola Normale Superiore, Pisa, Italy

GLEN MBENG

Department of Physics, Scuola Internazionale di Studi Superiori Avanzati, Istituto Nazionale di Fisica Nucleare, Trieste, Italy

AND

STEFANO SOATTO

Department of Computer Science, University of California, Los Angeles, CA, USA

[Received on 25 July 2018; revised on 9 July 2020; accepted on 5 October 2020]

We introduce an asymmetric distance in the space of learning tasks and a framework to compute their complexity. These concepts are foundational for the practice of transfer learning, whereby a parametric model is pre-trained for a task, and then fine tuned for another. The framework we develop is non-asymptotic, captures the finite nature of the training dataset and allows distinguishing learning from memorization. It encompasses, as special cases, classical notions from Kolmogorov complexity and Shannon and Fisher information. However, unlike some of those frameworks, it can be applied to large-scale models and real-world datasets. Our framework is the first to measure complexity in a way that accounts for the effect of the optimization scheme, which is critical in deep learning.

Keywords: task structure; task distance; information theory; Kolmogorov complexity; deep learning.

1. Introduction

The widespread use of deep learning is due in part to its flexibility: one can *pre-train* a deep neural network (DNN) for a task, say finding cats and dogs in images, and then *fine tune* it for another, say detecting tumors in a mammogram, or controlling a self-driving vehicle. Sometimes it works. So far, however, it has not been possible to predict whether such a ‘transfer learning’ practice will work, and how well. Even the most fundamental questions are still unanswered: how far are two tasks? In what space do learning tasks live? What is the complexity of a task? How difficult is it to transfer from one task to another? In this paper, we lay the foundations for answering these questions. The contributions of this paper can be summarized as follows:

1. We introduce a *distance between learning tasks* (Section 4), where tasks are represented by finite datasets of input data and discrete output classes (labels). Each task can have a different number of labels and a different number of samples. The distance behaves properly with respect to composition and inclusion relations (Corollary 4.1) and can handle corner cases, such as learning random labels, where other distances such as Kolmogorov’s fail (Theorem 3.4). The distance is

asymmetric by design, as it may be easier to learn one task starting from the solution to another, than the reverse (Theorem 4.1).

2. Some learning tasks are more difficult than others, so *complexity* plays a key role in their distance (Section 3). We define a new notion of complexity of a learning task that captures as special cases Kolmogorov complexity, Shannon mutual information and the Fisher information of the parameters of a model (Theorem 5.2, Theorem 5.3 and Theorem 5.4). It leverages and extends the classical notions of structure function and minimal sufficient statistics. It also provides a new interpretation to the PAC–Bayes generalization bound (Section 5.2).
3. We show how a parametric function class, such as DNNs trained with stochastic gradient descent (SGD), can be employed to measure the complexity of a learning task (Section 5). Such complexity is related to the ‘information in the parameters’ of the network. Experiments show that the resulting distance correlates with the ease of transfer learning by finetuning a pre-trained model (Section 7).
4. We show that the asymmetric distance between tasks is a lower bound to the cost of transfer learning using DNNs. We introduce the notion of *accessibility* (or *reachability*) of a learning task and illustrate examples of tasks that, while close to each other, are not accessible from one another. Accessibility depends on the global geometry of the space of tasks and on the optimization scheme employed. We conjecture that tasks are accessible through a local optimization scheme when there exists a continuous path of minimizers of the structure function (Section 6.2).

Our notion of complexity of a learning task formalizes intuitions regarding overfitting and generalization and differentiates *learning* from *memorization*. It is, to the best of our knowledge, the first complexity measure to take into account the training algorithm, rather than just its asymptotic outcome. It is also the first to enable practical computation for models with millions of parameters, as those currently in use in real applications.

1.1 Organization of the paper

The four main sections (3, 4, 5 and 6) are organized as a matrix: along the rows, Sections 3 and 4 use the language of Kolmogorov to introduce the notions of complexity, structure and distance of learning tasks, whereas Sections 5 and 6 introduce our more general framework, with classical concepts from Kolmogorov, Shannon and Fisher emerging as special cases. Along the columns, Sections 3 and 5 deal with *complexity* and the associated notion of information of a single task, whereas Sections 4 and 6 deal with the *distance* between tasks.

1.2 Related work

This work tackles the theoretical foundations of transfer learning, an active area of research too vast to review here (see [6] for a recent survey). More specifically, [18] computes pairwise distances between tasks by explicitly fine-tuning models for each pair. This method does not scale to a large number of tasks and only provides pairwise distances. On the other hand, [1] uses the Fisher information matrix to construct a linear embedding that enables to *predict* the effect of fine tuning without explicitly performing the computation. This provides empirical validation to our framework in Section 5. The existence of tasks that are similar according to most distances, yet one cannot ‘reach’ one from another,

has been observed empirically in [3]. This observation motivates our exploration of the dynamics of learning in Section 6. The notion of reachability in dynamical systems has been studied extensively for known models, whereas in our case the model changes as the training progresses. Numerical methods for reachability computation [12] are restricted to low-dimensional spaces, whereas typical models in use in the practice can have millions of parameters. Our work leverages classical results from Kolmogorov’s complexity theory [17], classical statistical inference and information theory and relates to recent theoretical frameworks for deep learning including the Information Bottleneck Principle [16], with important differences that we outline throughout the paper.

2. Preliminaries and nomenclature

In supervised learning, one is given a finite (training) dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1, \dots, N}$ of N samples, where $x_i \in \mathcal{X}$ is the input data (e.g., an image) and $y_i \in \mathcal{Y}$ is the output (e.g., a discrete label). The goal is to learn (i.e., estimate the parameters w of) a model p_w (a parametric function) that maps inputs x to estimated outputs \hat{y} , so that some loss (or risk) is minimized on unseen (test) data. It is common to assume that \mathcal{D} consists of i.i.d. samples from some unknown distribution $p(x, y)$, which are used to compute a loss function that is minimized with respect to the parameters w so that $p_w(y|x)$ is close to the ‘true’ posterior $p(y|x)$.

In this work, however, we make no assumption on the data generation process, nor do we attempt to approximate the true posterior $p(y|x)$, even if one existed. Instead, we adopt Kolmogorov’s approach to learning the ‘structure’ in the data, directly accounting for the finite sample size N , the approximation error in the loss $L_{\mathcal{D}}$ and the complexity of the model. The basic elements of Kolmogorov complexity theory that we use are described in [17] and summarized in Section 5.1 to make the paper self-contained.

In deep learning, the optimization scheme acts as an implicit regularizer in the loss [5]. Hence, for a measure of complexity to be relevant to deep learning it must take into account both the loss function and the optimization scheme. To the best of our knowledge, none of the classical measures do so.

2.1 Deep neural networks

DNNs are a class of parametrized functions (models) obtained by successive composition of ‘layers’ of affine operations, where both the linear terms and offset parameters are referred to as *weights*, followed by element-wise non-linear operations such as a saturation or rectification. One of the most common non-linearities is the rectified linear unit (ReLU), which leaves the positive part unchanged and zeroes the negative part. The first layer is thus of the form $x \mapsto h(Wx + b)$, where (W, b) are the weights, or improperly the ‘weight vector’; $h(x)$ is defined as the component-wise maximum between 0 and x . The output of the first layer is therefore given by $z = h(Wx + b)$. The second layer performs an operation of the same type, with a different set of weights, taking as input the output of the first layer, and so on. The last layer produces a normalized (probability) vector $p_w(y|x)$, with components $y \in \{1, \dots, K\}$, usually through a soft-max operation, i.e., $z \mapsto e^z / \sum_k e^{z_k}$. The k -th entry of this vector represents, ideally, the probability $p_w(y = k|x)$ of the input x being of class $y = k$, as assessed by the model. The number of weights is typically larger than the number of samples in the training set.

The learning criterion is given by maximum likelihood: $\hat{w} = \arg \max_w L_{\mathcal{D}}(p_w)$, where $L_{\mathcal{D}}(p_w)$ is the *loss function* and is given by $L_{\mathcal{D}}(p_w) = \sum_{i=1}^N -\log p_w(y_i|x_i)$. This loss can be interpreted as an empirical approximation of the average cross-entropy $H_{p_w, p_w}(y|x)$, which is zero when $p_w(y|x) = p(y|x)$. Note, however, that the latter posterior probability $p(y|x)$ is *not* among the infinitely many minimizers

of the empirical loss,¹ which therefore needs to be regularized. In deep learning, regularization is both explicit and implicit. Explicit mechanisms include additional terms in the loss function such as the norm of the weights (weight decay), or characteristics of the function class (architecture) such as downsampling and averaging (pooling). Implicit mechanisms include characteristics of the optimization.

The training process consists of minimizing the loss using one of the many variants of SGD. At every iteration, SGD makes a step in the direction of the (negative) gradient of $L_{\mathcal{D}}$, which is approximated using a random subset of the training set (*minibatch*). The length of the step is a hyperparameter called the *learning rate*.

A learning task can be *described* by the relation between the input variable x (e.g., an image) and the corresponding output variable y (e.g., ‘cat’ or ‘dog’) and would be *defined* by the posterior distribution $p(y|x)$. Since that is typically not available, in practice the task is *specified* by a dataset \mathcal{D} (e.g., a collection of images with a label of either ‘cat’ or ‘dog’): the dataset is all that is known about the desired goal before the training process starts. The *solution* to the task is determined by the dataset along with various *biases*, corresponding to modeling choices of loss function, class of functions and optimization method.

3. Complexity of a learning task

The same task could be specified by different datasets, with different numbers of samples and different size of the inputs. So, the complexity of the task is not just the complexity of the dataset: the popular MNIST and CIFAR-10 classification tasks are similar on these counts, yet one is very ‘simple’ to learn, the other is not. Instead, the complexity of a dataset is a function of two factors: the underlying *structure* that describes the data collectively, as well as *variability* that individual datapoints exhibit relative to the shared structure. This split is not unique: a dataset can be described or explained with different structures, each leaving a different amount of residual variability unexplained. These two factors are captured in the following definition.

DEFINITION 3.1 Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1, \dots, N} \subseteq \mathcal{X} \times \mathcal{Y}$ be a dataset. We define the complexity of \mathcal{D} as

$$C(\mathcal{D}) = \min_{p(y|x)} L_{\mathcal{D}}(p) + K(p), \quad (3.1)$$

where $L_{\mathcal{D}}(p) = \sum_{i=1}^N -\log p(y_i|x_i)$ is the empirical classification (cross-entropy) loss, and the minimum is over all possible computable probability distributions $p(y|x)$ of the label y given the input x . By $K(p)$ we denote the Kolmogorov complexity of the distribution $p(y|x)$.

At first sight, $C(\mathcal{D})$ appears similar to a conditioned version of the two-part code in [17, Appendix II]:

$$C_K(\mathcal{D}) = \min_{p(y|\mathbf{x})} -\log p(y|\mathbf{x}) + K(p), \quad (3.2)$$

where $\mathbf{y} = \langle y_1, \dots, y_N \rangle$ and $\mathbf{x} = \langle x_1, \dots, x_N \rangle$ are strings obtained by concatenating all labels and inputs of \mathcal{D} , respectively. In Theorem 3.2.1, we recall that $C_K(\mathcal{D})$ coincides with the conditional Kolmogorov

¹ The minimizers are of the form $\sum_{i=1}^N \delta(y_i - \phi(x_i))$ for any function ϕ such that $\phi(x_i) = y_i$. Such minimizers are said to *overfit* the data.

complexity $K(y|x)$ of the string y given x . Instead, in eq. 3.1, we only consider factorized distributions $p(y|x) = \prod_i p(y_i|x_i)$. This has major consequences for determining the complexity of a task: the distribution $p(y|x)$ minimizing eq. 3.2 does not need to encode in p all task-relevant information, as it can, at test time, extract any missing information from the training set x (cf. Theorem 3.2.3). Hence, $K(p)$ in eq. 3.2 alone would not be a valid measure of complexity. Instead, the distribution $p(y|x)$ minimizing eq. 3.1 can only access a single datum x at test time; hence, all structure in the dataset has to be encoded in p . In other words, the solution of eq. 3.2 does not need to capture the structure in D , and therefore its complexity is not reflective of the complexity of the task. On the other hand, the solution of eq. 3.1 is forced to capture the structure in the data.

Also note that, unlike $C_K(\mathcal{D})$, $C(\mathcal{D})$ is invariant to permutations of the dataset. Suppose that a dataset \mathcal{D} for a binary classification task is ordered so that all negative samples \mathcal{D}_- precede the positive samples \mathcal{D}_+ . Then we would have $C_K(\mathcal{D}) \leq \log |\mathcal{D}_-| + \log |\mathcal{D}_+|$ regardless of the complexity of the task, as it suffices to encode the number of negative and positive samples to reproduce the string y exactly. However, we show in Theorem 3.2.3 that making $p(y|x)$ permutation invariant does not yield a sensible measure of complexity of a task. To address permutation invariance, [11] proposes the following definition, which uses deterministic functions rather than probability distributions:

$$C_{\text{DET}}(\mathcal{D}) = \min \{K(f) \mid f: \mathcal{X} \rightarrow \mathcal{Y} \text{ such that } f(x_i) = y_i \ \forall (x_i, y_i) \in \mathcal{D}\}.$$

The following proposition compares these definitions of complexity.

PROPOSITION 3.2 (Measures of complexity). Up to an additive term that does not depend on the dataset \mathcal{D} , we have the following:

1. $K(y|x) = C_K(\mathcal{D})$.
2. $C_K(\pi(\mathcal{D})) \leq C(\mathcal{D})$ for any permutation $\pi(\mathcal{D})$ of \mathcal{D} .
3. For every $C > 0$, there is a dataset \mathcal{D} such that $C(\mathcal{D}) \geq C$ and $C_K(\pi(\mathcal{D})) = O(1)$ for any permutation $\pi(\mathcal{D})$. Therefore, $C(\mathcal{D})$ is not simply the (average) complexity $C_K(\pi(\mathcal{D}))$ of encoding some permutation of \mathcal{D} .
4. When $C_{\text{DET}}(\mathcal{D})$ is defined, i.e., if there is a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ such that $f(x_i) = y_i$ for each $(x_i, y_i) \in \mathcal{D}$, then we have $C(\mathcal{D}) \leq C_{\text{DET}}(\mathcal{D})$. In addition, $C(\mathcal{D}) = C_{\text{DET}}(\mathcal{D})$ if we have an oracle that provides a bijective map $\{x_1, \dots, x_N\} \rightarrow \{1, \dots, N\}$.

Assuming the data are sampled i.i.d. from a computable probability distribution p , the following proposition characterizes the complexity of the dataset. In particular, it shows that, asymptotically, the complexity of the dataset is given by the noise in the labels and the complexity of the distribution generating the data. How to ignore the effect of noise on the complexity, and what happens in the non-asymptotic regime, is central to Kolmogorov's framework, which we will build upon in the next sections.

PROPOSITION 3.3 Fix a probability distribution $p(x, y)$ on $\mathcal{X} \times \mathcal{Y}$, and assume that $p(y|x)$ is computable. If \mathcal{D} is a collection of N i.i.d. samples $(x_i, y_i) \sim p(x, y)$, then:

1. The expected value of $C(\mathcal{D})$ satisfies

$$N \cdot H_p(y|x) \leq \mathbb{E}[C(\mathcal{D})] \leq N \cdot H_p(y|x) + K(p(y|x)),$$

where $H_p(y|x)$ is the conditional entropy of y given x .

2. For any $\epsilon > 0$ there is N_0 such that, with probability $1 - \epsilon$, for any $N \geq N_0$ we have the equality

$$C(\mathcal{D}) = N \cdot H_p(y|x) + K(p(y|x)),$$

and p is the only computable distribution for which the equality holds.

It is instructive to test our definition of complexity on a dataset of random labels, whereby each input is assigned a label at random. We will revisit this case often, as it challenges many of the extant theories of deep learning [19].

EXAMPLE 3.4 (Complexity of random labels). Suppose that each input x_i of the dataset \mathcal{D} is associated to a label $y_i \in \mathcal{Y}$ sampled uniformly at random in a fixed finite set, so $p(y|x) = 1/|\mathcal{Y}|$ has a constant complexity. Under the same assumptions of Theorem 3.2.4, the expected value of $C(\mathcal{D})$ is $N \log |\mathcal{Y}| + O(1)$ by Theorem 3.3. Since $C(\mathcal{D}) \leq N \log |\mathcal{Y}| + O(1)$ for any such \mathcal{D} , the complexity of a ‘typical’ dataset with random labels is approximately $N \log |\mathcal{Y}|$.

In a sense, learning random labels is a very complex task: they cannot be predicted from the input, so the model is forced to *memorize* them (‘overfitting’). Accordingly, $C(\mathcal{D})$ in Theorem 3.4 is very high. However, by construction, there is *no structure in the data*, which is why the model cannot generalize to unseen data. In other words, this is a complex memorization task, but a *trivial learning task*, as there is nothing to learn from the data. We would like the definition of complexity to reflect this distinction between learning and memorization.

Another important aspect not captured by this definition of complexity is the role of performance in learning the structure in the data. For example, one can train a trivial model to distinguish images of airplanes from fireplaces by counting the number of blue pixels. It will not be very precise. To achieve a small error, however, one must learn what makes an airplane different from a fireplace even if the latter is painted blue.

3.1 Structure function of a task

The trade-off between the loss achievable by a solution p on a dataset \mathcal{D} and its complexity is captured by the *structure function* [17]:

$$S_{\mathcal{D}}(t) = \min_{K(p) \leq t} L_{\mathcal{D}}(p). \quad (3.3)$$

It is a decreasing function that reaches zero for sufficiently high complexity, depending on the task. As we increase complexity, the loss decreases rapidly while simple models correctly classify easy samples. After all shared structure is captured, the only way to further reduce the loss is to memorize the leftover samples, leading to the worst possible trade-off of one nat of complexity for each unit of decrease in loss. Eventually, training enters this *linear (overfitting) regime for every dataset, given sufficient complexity*. For random labels, the entire training is in this overfitting regime since there is no structure in the data.

EXAMPLE 3.5 (Structure function for random labels). By definition of $C(\mathcal{D})$ we have $C(\mathcal{D}) \leq L_{\mathcal{D}}(p) + K(p)$ for any computable probability distribution $p(y|x)$. For a typical dataset with random labels (see Theorem 3.4), we have that $C(\mathcal{D}) \simeq N \log |\mathcal{Y}|$. Therefore

$$L_{\mathcal{D}}(p) \gtrsim N \log |\mathcal{Y}| - K(p), \quad \text{and hence} \quad S_{\mathcal{D}}(t) \gtrsim N \log |\mathcal{Y}| - t.$$

The lower bound for $S_{\mathcal{D}}(t)$ can be achieved by memorizing the label of $\lfloor t/\log |\mathcal{Y}| \rfloor$ data points. Therefore $S_{\mathcal{D}}(t) \simeq N \log |\mathcal{Y}| - t$.

The structure function of a dataset cannot be computed in general ([17, Section VII]). In Section 5 we introduce a generalized version that is computable: Figure 2 shows the result on common datasets. The predicted fast decrease in the loss as complexity increases is followed by the asymptotic linear phase. The sharp transition to the linear regime is clearly visible as a function of the loss: as the parameter β weighting complexity increases, a plateau is reached that depends on the task. Note that for random labels the loss decreases linearly as expected (left).

3.2 Task Lagrangian and minimal sufficiency

The constrained optimization problem in the definition of the structure function (eq. 3.3) has an associated Lagrangian $L_{\mathcal{D}}(p) + \beta K(p)$, where β is a Lagrange multiplier that trades off the complexity of the model $K(p)$ with the fidelity $L_{\mathcal{D}}(p)$. If we take the minimum over p , we obtain a family of complexity measures parametrized by β :

$$C_{\beta}(\mathcal{D}) = \min_{p(y|x)} L_{\mathcal{D}}(p) + \beta K(p). \quad (3.4)$$

As a function of β , this is the Legendre transform of the structure function $S(t)$. To minimize $L_{\mathcal{D}}(p) + \beta K(p)$, we can increase the complexity $K(p)$ of the model until the return of doing so has a ratio that is smaller than the constant β we have selected.

If p^* is a minimizer of eq. 3.4 for $\beta = 1$, the corresponding Kolmogorov complexity $t^* = K(p^*)$ is the value at which the structure function $S(t)$ reaches a linear regime. Thus, the special case $\beta = 1$ marks the transition to overfitting and is related to Kolmogorov's notion of minimal sufficient statistic [17]. Since we are using Kolmogorov's complexity, $\beta = 1$ is the worst possible trade-off.

PROPOSITION 3.6 Given a task \mathcal{D} , let β^* be the largest β for which $C_{\beta}(\mathcal{D})$ is not realized by a constant distribution $p(y|x) = p(y)$. Then $\beta^* \geq 1$, and $\beta^* = 1$ if \mathcal{D} is a typical dataset with random labels.

As we have seen for random labels, a dataset may be complex and yet exhibit little underlying structure. We say that a distribution $p(y|x)$ is a Kolmogorov *sufficient statistic* of \mathcal{D} if it minimizes $C(\mathcal{D})$. It is *minimal* if it also minimizes $K(p)$ among all sufficient statistics, that is, the smallest statistic that is able to solve the task optimally. The rationale is that the smallest statistic that solves a task should not squander resources by memorizing nuisance variability. Rather, it should only capture the important information shared among the data. This is shown in the following example.

EXAMPLE 3.7 For random labels, both the distribution $p(y|x)$ that memorizes all the labels in the dataset, and the uniform distribution $p(y|x) = 1/|\mathcal{Y}|$, are sufficient statistics. However, only the latter is minimal, since $K(p)$ is a constant that does not depend on \mathcal{D} . There is no structure to be extracted from a dataset of random labels.

The level of complexity of a model is an important design parameter the practitioner wishes to control. Rather than seeking minimal sufficient statistics, we explore the entire trade space, by introducing the notion of β -sufficiency.

DEFINITION 3.8 Given a dataset \mathcal{D} , define a β -sufficient statistic of \mathcal{D} as a probability distribution $p(y|x)$ such that $C_{\beta}(\mathcal{D}) = L_{\mathcal{D}}(p) + \beta K(p)$. We say that $p(y|x)$ is a β -minimal sufficient statistic if it also minimizes $K(p)$ among all β -sufficient statistics.

Notice that, for $\beta = 1$, Theorem 3.8 reduces to minimal sufficiency in the sense of Kolmogorov.

4. Asymmetric distance between tasks

We now introduce a one-parameter family of distances between tasks. The hope is for them to correlate with the success of transfer learning. Since it is easier to learn a simple task from a complex one, it is desirable for the distance to be asymmetric.

DEFINITION 4.1 The distance from task \mathcal{D}_1 to \mathcal{D}_2 at level β is

$$d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2) := \max_{p_1} \min_{p_2} K(p_2|p_1) = \max_{p_1} \min_{p_2} K(\langle p_1, p_2 \rangle) - K(p_1),$$

where p_i varies among all β -minimal sufficient statistics of \mathcal{D}_i .

The intuition behind this definition is the following: for a task \mathcal{D}_2 to be close to \mathcal{D}_1 , every β -minimal sufficient statistic p_1 of \mathcal{D}_1 should be close to some β -minimal sufficient statistic p_2 of \mathcal{D}_2 . Then, every optimal model p_1 of \mathcal{D}_1 can be fine tuned to some optimal model p_2 of \mathcal{D}_2 .

LEMMA 4.1 The asymmetric distance d_β satisfies the following properties:

- $d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2) \geq 0$ (positivity);
- $d_\beta(\mathcal{D} \rightarrow \mathcal{D}) = O(1)$ (a task is close to itself);
- $d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_3) \leq d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2) + d_\beta(\mathcal{D}_2 \rightarrow \mathcal{D}_3) + O(1)$ (triangle inequality).

We now derive a characterization of the distance between tasks based on the complexity of their composition, amenable to generalization in Section 6. Denote by $\mathcal{D}_1 \sqcup \mathcal{D}_2$ the disjoint union of two datasets \mathcal{D}_1 and \mathcal{D}_2 , defined as

$$\mathcal{D}_1 \sqcup \mathcal{D}_2 = \{(i, x), y \mid i \in \{1, 2\}, (x, y) \in \mathcal{D}_i\}.$$

Notice that an index i is added to the input, in order to recognize the original dataset. A desirable property for a distance between tasks would be that $d_\beta(\mathcal{D}_1 \sqcup \mathcal{D}_2 \rightarrow \mathcal{D}_1) = O(1)$: indeed, a model that performs well on $\mathcal{D}_1 \sqcup \mathcal{D}_2$ should be easily fine tuned to a model that performs well on \mathcal{D}_1 alone. Adding the index i in the definition of $\mathcal{D}_1 \sqcup \mathcal{D}_2$ is essential, as we can see in the following example.

EXAMPLE 4.2 Let \mathcal{D} be a typical dataset with random labels, and let $\mathcal{D}_1 \subseteq \mathcal{D}$ be the set of data points $(x, y) \in \mathcal{D}$ satisfying some property of Kolmogorov complexity $t \gg 0$. Then the whole dataset \mathcal{D} has a trivial structure, whereas \mathcal{D}_1 has a complicated structure. If both $|\mathcal{D}_1|$ and $|\mathcal{D} \setminus \mathcal{D}_1|$ are large, by Theorem 3.3, the Kolmogorov complexity of a minimal sufficient statistic of \mathcal{D}_1 is t , and the complexity of a minimal sufficient statistic of \mathcal{D} is $O(1)$.

We now prove that, under the hypotheses of Theorem 3.3.2, the property $d_\beta(\mathcal{D}_1 \sqcup \mathcal{D}_2 \rightarrow \mathcal{D}_1) = O(1)$ is satisfied.

THEOREM 4.3 Suppose that \mathcal{D}_1 and \mathcal{D}_2 are obtained by sampling from two fixed distributions p_1 and p_2 on $\mathcal{X} \times \mathcal{Y}$, such that $p_1(y|x)$ and $p_2(y|x)$ are computable. Then, with high probability and for $|\mathcal{D}_1|$ and $|\mathcal{D}_2|$ sufficiently large, $d_\beta(\mathcal{D}_1 \sqcup \mathcal{D}_2 \rightarrow \mathcal{D}_1) = O(1)$.

COROLLARY 4.1 Under the hypothesis of Theorem 4.3, we have that

$$d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2) = d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_1 \sqcup \mathcal{D}_2) + O(1) = \max_{p_1} \min_{p_{12}} K(p_{12}) - K(p_1) + O(1),$$

where p_{12} varies among the β -minimal sufficient statistics of \mathcal{D}_{12} and p_{12} varies among those of \mathcal{D}_1 .

We now have a way of comparing different learning tasks, at least in theory. The asymmetric distance $d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2)$ allows us to quantify how difficult it is to learn a new task \mathcal{D}_2 given a solution to \mathcal{D}_1 . However, quantities defined in terms of Kolmogorov complexity are difficult to handle in practice and may behave well only in an asymptotic regime. In the next section, we introduce a generalization of the framework developed so far that can be instantiated for a particular model class such as DNNs.

5. Information in the model parameters

Whatever structure or ‘information’ was captured from the dataset, it ought to be measurable from the model parameters, since they are all we have left after training. As we will see in the next section, this intuition is faulty, as *how* we converge to a set of parameter (i.e., the optimization algorithm) also affects what information we can extract from the data. For now, we focus on generalizing the theory in the previous section with an eye towards computability. Although most of our arguments are general, we focus on DNNs as a model class. They can have millions of parameters, so measuring their information can be non-trivial.

One way to compute information in the parameters is to measure their coding length at some level of precision, independent of the particular task. This is suboptimal, as only a small subset of the weights of a trained neural networks matters: imagine changing a certain component of the weights, and observing no change in the loss. Arguably, that weight ‘contains no information’ about the dataset. For the purpose of storing the trained model, that weight could be replaced with any constant, or randomized each time the network is used. The loss landscape has a small curvature² in the coordinate direction corresponding to that weight. On the other hand, imagine changing the least significant bit of another component of the weights and noticing a large increase in the loss. That weight could be said to be very ‘informative,’ and its valued should be stored with high precision.

With these observations in mind, we allow the weights to be encoded with some uncertainty, through a probability distribution $Q(w|\mathcal{D})$, which depends on the dataset \mathcal{D} . For example, Dirac’s Delta $Q(w|\mathcal{D}) = \delta_{w^*}$ corresponds to an exact encoding of the weight vector w^* . If we fix a reference ‘pre-distribution’ $P(w)$, [10] shows that the labels y can be reconstructed from the input x and the pre-distribution $P(w)$, by using

$$\mathbb{E}_{w \sim Q(w|\mathcal{D})} [L_{\mathcal{D}}(p_w(y|x))] + \text{KL}(Q(w|\mathcal{D}) \parallel P(w))$$

additional nats. This expression resembles the right-hand side of eq. 3.1 in that it describes a trade-off between fidelity and complexity. Here, Kolmogorov complexity has been replaced by the Kullback–Liebler (KL) divergence $\text{KL}(Q(w|\mathcal{D}) \parallel P(w))$, which we call the *information in the parameters* of the

² We will elaborate on this point after Theorem 5.4.

model.³ This leads to the following new definition of complexity. It entails two arbitrary choices of distributions: one that is chosen before any data are observed, which we call ‘pre-distribution’ $P(w)$, and one that is chosen at the end of training, which we call ‘post-distribution’ $Q(w|D)$. We deliberately avoid calling them ‘prior’ and ‘posterior’ to avoid any confusion with a Bayesian setting. Here, the weights w are fixed at the end of training, so if the post-distribution was the true posterior, it would be degenerate. One should also not confuse the post-distribution with the distribution of the weights *during training*, due for instance to the stochasticity of the optimization algorithm, on which we will comment later.

DEFINITION 5.1 The complexity of the task \mathcal{D} at level β , using the post-distribution $Q(w|\mathcal{D})$ and the pre-distribution $P(w)$, is given by

$$C_\beta(\mathcal{D}; P, Q) = \mathbb{E}_{w \sim Q(w|\mathcal{D})} [L_{\mathcal{D}}(p_w(y|x))] + \beta \underbrace{\text{KL}(Q(w|\mathcal{D}) \parallel P(w))}_{\text{information in the parameters}}. \quad (5.1)$$

The second term, $\text{KL}(Q(w|\mathcal{D}) \parallel P(w))$, measures the information in the parameters of the model. We refer to $\mathbb{E}_{w \sim Q(w|\mathcal{D})} [L_{\mathcal{D}}(p_w(y|x))]$ as the (expected) reconstruction error of the label under the hypothesis $Q(w|\mathcal{D})$.

We emphasize again that there is no implied Bayesian interpretation, as $Q(w|\mathcal{D})$ can be *any* distribution. Depending on the choice, this expression can be computed in closed form or estimated (Section 5.1). For instance, when $Q(w|\mathcal{D}) = \delta_{w^*}$, the expression reduces to the length of a two-part code for \mathcal{D} using the model class $p_w(y|x)$. However, Theorem 5.1 is more general and can be extended to the continuous case, or in cases where there is a *bona fide* distribution, as in variational inference and Bayesian neural networks.

Another fundamental difference of our framework compared to prior work in the area is that, while eq. 3.1 measures the complexity in terms of the best obtainable by the model class (in that case, the class of computable probability distributions), the complexity $C_\beta(\mathcal{D}; P, Q)$ takes into account both the particular model class and the *training algorithm*, i.e., the map $A : \mathcal{D} \mapsto Q(w|\mathcal{D})$, as we shall see in Section 6.2.

5.1 Relation with Kolmogorov, Shannon and Fisher information

Of all the possible choices of pre-distribution $P(w)$ in Theorem 5.1, we investigate three special cases: The ‘universal prior’ of all computable distributions; an ‘adapted prior’ that relies on a probability distribution over datasets; an uninformative prior, agnostic of the dataset.

We start with the first case, which provides a link between Theorem 5.1 and the framework of Section 3. For a given weight vector w , we define the universal prior $P(w) = \frac{1}{Z} e^{-K(w)}$, where Z is a normalization constant. This can be interpreted as follows: for every w , choose a minimal program that outputs w , and assign it a probability that decreases exponentially in terms of the length of the program.

³ This should not be confused with the information in x and y , or any intermediate representation that is build by the model, such as the activations of a DNN, which is more frequently studied. Information in the parameters and information in the activations are different and, in the case of DNNs, are related through the emergence bound [4].

PROPOSITION 5.2 (Kolmogorov complexity of the weights). Let $P(w)$ be the universal prior, and let $Q(w|\mathcal{D}) = \delta_{w^*}$ be a Dirac delta. Then the information in the weights equals the Kolmogorov complexity of the weights $K(w^*)$, up to a constant.

We now turn to the second case, which provides a link with Shannon mutual information.

PROPOSITION 5.3 (Shannon information in the weights). Assume the dataset \mathcal{D} is sampled from a distribution $\pi(\mathcal{D})$, and let the outcome of training on a sampled dataset \mathcal{D} be described by a distribution $Q(w|\mathcal{D})$. Then the pre-distribution $P(w)$ minimizing the expected complexity $\mathbb{E}_{\mathcal{D}}[C_{\beta}(P, Q)]$ is $P(w) = \mathbb{E}_{\mathcal{D}}[Q(w|\mathcal{D})]$, and the expected information in the weights is given by

$$\mathbb{E}_{\mathcal{D}}[\text{KL}(Q(w|\mathcal{D}) \| P(w))] = I(w; \mathcal{D}).$$

Here $I(w; \mathcal{D})$ is Shannon's mutual information between the weights and the dataset, where the weights are seen as a (stochastic) function of the dataset given by the training algorithm (e.g., SGD).

Note that, in this case, the pre-distribution $P(w)$ is optimal given the choice of the training algorithm (i.e., the map $A : \mathcal{D} \rightarrow Q(w|\mathcal{D})$) and the distribution of training datasets $\pi(\mathcal{D})$. However, the distribution $\pi(\mathcal{D})$ is generally unknown, as we are often given a single dataset to train. Even if it was known, computing the marginal distribution $\mathbb{E}_{\mathcal{D}}[Q(w|\mathcal{D})]$ over all possible datasets would not be realistic, as it is high dimensional and has complex interactions between different components. Nevertheless, it is interesting that the information in the parameters specializes to Shannon's mutual information in the weights [4].

The third case, namely an uninformative prior with a Gaussian posterior, is the most practical and provides a link to the Fisher information matrix and the learning dynamics of common optimization algorithms such as SGD.

THEOREM 5.4 (Fisher information in the weights). Choose an isotropic Gaussian pre-distribution $P(w) \sim N(0, \lambda^2 I)$. Let the post-distribution $Q(w|\mathcal{D})$ also be Gaussian: $Q(w|\mathcal{D}) \sim N(w^*, \Sigma)$, where w^* is a local minimum of the cross-entropy loss. Then, for $\lambda \rightarrow \infty$, we have that:

- the covariance Σ^* that minimizes $C_{\beta}(P, Q)$ tends to $\beta H^{-1} = \frac{\beta}{N} F^{-1}$ (this is in accordance with the Cramér–Rao bound);
- the information in the weights is given by

$$\text{KL}(Q(w|\mathcal{D}) \| P(w)) = \frac{1}{2} \log |F| + \frac{1}{2} k \log \lambda^2 + O(1).$$

Recalling that the Fisher information F measures the local curvature, this proposition confirms the qualitative discussion at the beginning of Section 5: the optimal covariance $\Sigma^* \propto H^{-1}$ gives high variability to the directions of low curvature, which are ‘less informative’, whereas it gives low variability to the ‘more informative’ directions of high curvature. The Fisher information describes the information contained in the weights about the dataset. In Section 6 we discuss how to compute it.

5.2 Connections with the PAC–Bayes bound

The Lagrangian $C_{\beta}(P, Q)$ admits another interpretation as an upper-bound to the test error, as shown by the PAC–Bayes test error bound:

THEOREM 5.5 ([14, Theorem 2]). Assume the dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ is sampled i.i.d. from a distribution $p(y, x)$, and assume that the per-sample loss used to train is bounded by $L_{\max} = 1$ (we can reduce to this case by clipping and rescaling the loss). For any fixed $\beta > 1/2$, pre-distribution $P(w)$ and post-distribution $Q(w|\mathcal{D})$, with probability at least $1 - \delta$ over the sample of \mathcal{D} , we have

$$\begin{aligned} L_{\text{test}}(Q) &\leq \frac{1}{N(1 - \frac{1}{2\beta})} \left[\mathbb{E}_{w \sim Q(w|\mathcal{D})} [L_{\mathcal{D}}(p_w)] + \beta \text{KL}(Q \| P) + \beta \log \frac{1}{\delta} \right] \\ &= \frac{1}{N(1 - \frac{1}{2\beta})} \left[C_{\beta}(\mathcal{D}; P, Q) + \beta \log \frac{1}{\delta} \right], \end{aligned}$$

where $L_{\text{test}}(Q) := \mathbb{E}_{x, y \sim p(x, y)} [\mathbb{E}_{w \sim Q} [p_w(y|x)]]$ is the expected per-sample test error that the model incurs using the weight distribution $Q(w|\mathcal{D})$.

Hence, we see that minimizing the Lagrangian $C_{\beta}(\mathcal{D}; P, Q)$ can be interpreted as minimizing an upper-bound on the test error of the model, rather than directly minimizing the train error. This is in accordance with the intuition developed earlier that minimizing $C_{\beta}(\mathcal{D}; P, Q)$ forces the model to capture the structure of the data. It is also interesting to consider the following bound on the expectation over the sampling of \mathcal{D} ([14, Theorem 4]):

$$\mathbb{E}_{\mathcal{D}} [L_{\text{test}}(Q(w|\mathcal{D}))] \leq \frac{1}{N(1 - \frac{1}{2\beta})} \left[\mathbb{E}_{\mathcal{D}} [L_{\mathcal{D}}(Q(w|\mathcal{D}))] + \beta \mathbb{E}_{\mathcal{D}} [\text{KL}(Q(w|\mathcal{D}) \| P)] \right].$$

As we have seen in Theorem 5.3, for the optimal choice of pre-distribution P minimizing the bound, we have $\mathbb{E}_{\mathcal{D}} [\text{KL}(Q \| P)] = I(w; \mathcal{D})$. Hence, the Shannon information that the weights of the model have about the dataset \mathcal{D} is the measure of complexity that gives (on expectation) the strongest generalization bound. This has also been noted in [4]. In [8], a non-vacuous generalization bound is computed for DNNs, using (non-centered and non-isotropic) Gaussian prior and posterior distributions.

6. Generalized distance, reachability and learnability of tasks

Unlike $C_{\beta}(\mathcal{D})$, the definition of $C_{\beta}(\mathcal{D}; P, Q)$ in eq. 5.1 captures the complexity of a dataset for a particular model class and training algorithm. Motivated by this, we now define a distance between datasets, which is tailored to the model.

Throughout this section, fix a parametrized model class $p_w(y|x)$, a pre-distribution $P(w)$ and a class \mathcal{Q} of post-distributions $Q(w|\mathcal{D})$. The case we are most interested in is that of DNNs, with an uninformative pre- and a Gaussian post-distribution (see Section 5.1). Our starting point is to generalize Kolmogorov's structure function. Consider the following generalized structure function:

$$S_{\mathcal{D}}(t) = \min_{\text{KL}(Q(w|\mathcal{D}) \| P(w)) \leq t} \mathbb{E}_{w \sim Q(w|\mathcal{D})} [L_{\mathcal{D}}(p_w)].$$

Here the minimum is taken among all post-distributions $Q(w|\mathcal{D})$ in the chosen class \mathcal{Q} . Similarly to what we have seen in Section 3.2, this minimization problem has $C_{\beta}(\mathcal{D}; P, Q)$ as its associated Lagrangian. We say that $Q(w|\mathcal{D})$ is a β -sufficient statistic if it minimizes $C_{\beta}(\mathcal{D}; P, Q)$. It is a β -minimal sufficient

statistic if it also minimizes $\text{KL}(Q(w|\mathcal{D}) \| P(w))$. Motivated by Corollary 4.1, we then introduce the following distance.

DEFINITION 6.1 The task distance from \mathcal{D}_1 to \mathcal{D}_2 at level β is

$$d_\beta(\mathcal{D}_1 \rightarrow \mathcal{D}_2) = \max_{Q_1} \min_{Q_{12}} \text{KL}(Q_{12}(w|\mathcal{D}) \| P(w)) - \text{KL}(Q_1(w|\mathcal{D}) \| P(w)),$$

where $Q_{12}(w|\mathcal{D})$ is a β -minimal sufficient statistic for $\mathcal{D}_1 \sqcup \mathcal{D}_2$, and $Q_1(w|\mathcal{D})$ is a β -minimal sufficient statistic for \mathcal{D}_1 .

While more general and amenable to computation than the distance of Theorem 4.1, this is not an actual distance as it lacks several of the properties in Lemma 4.1. Nonetheless, we will show that d_β does indeed capture the difficulty of fine-tuning from one dataset to another using specific model families (such as DNNs) and training algorithms (SGD) and empirically shows good correlation with other distances (e.g., taxonomical), when those are defined.

6.1 Local task reachability

Until now, we have only considered global minimization problems, where we aim to find the best solution satisfying some optimal trade-off. However, many learning algorithms (e.g., SGD) are local: starting from the current solution, they take a small step to greedily minimize some objective function. This raises the question of which conditions allow such an algorithm to recover an optimal solution.

Given a distribution $Q(w|\mathcal{D}) \in \mathcal{Q}$, denote by $L_{\mathcal{Q}}(Q)$ the expected loss $\mathbb{E}_{w \sim Q(w|\mathcal{D})}[L_{\mathcal{Q}}(p_w)]$. Fix a distance d on \mathcal{Q} such that both $L_{\mathcal{Q}}(Q)$ and $\text{KL}(Q \| P)$ are continuous, as functions $\mathcal{Q} \rightarrow \mathbb{R}$. This way, the Lagrangian $C_\beta(\mathcal{D}; P, Q) = L_{\mathcal{Q}}(Q) + \beta \text{KL}(Q \| P)$ is continuous in the joint variable (Q, β) . In the case of DNNs, where P is uninformative and \mathcal{Q} is the class of the Gaussian distributions $Q(w|\mathcal{D}) \sim N(w^*, \Sigma)$, we can take for example d as the Wasserstein distance, or the Euclidean distance between the parameters (w^*, Σ) .

DEFINITION 6.2 (ϵ -local learning algorithm). Fix $\beta \geq 0$. We say that a step is ϵ -local if, starting from a given statistic Q_0 , it finds the statistic Q that minimizes

$$C_\beta(\mathcal{D}; P, Q) = L_{\mathcal{Q}}(Q) + \beta \text{KL}(Q \| P)$$

and such that $d(Q, Q_0) \leq \epsilon$. We say that a learning algorithm is ϵ -local if it only takes ϵ -local steps.

In the limit $\epsilon \rightarrow 0$, this reduces to gradient descent on the Lagrangian C_β . Notice however that this is not the same as performing gradient descent on the cross-entropy loss $L_{\mathcal{Q}}$, unless $\beta = 0$. Indeed, minimizing $L_{\mathcal{Q}}$, unlike minimizing C_β (Section 5.2), gives no guarantees on the performance on test data, as the learning algorithm could simply memorize the dataset \mathcal{D} .

We will show in the next section that a DNN trained with SGD can actually be interpreted as a local learning algorithm minimizing C_β . Common methods of training DNNs also rely on annealing of the learning rate: training starts with a high learning rate that anneals gradually. As we will make precise in the next section, this corresponds to starting with a high value of β , and gradually decreasing it to a final value $\bar{\beta}$. This helps avoiding degenerate solutions, because the model starts by favouring structure over memorization. We introduce the following definition to capture the role of annealing.

DEFINITION 6.3 (ϵ -local learning algorithm with annealing). An ϵ -local learning algorithm with annealing is a learning algorithm that alternates ϵ -local steps (that change the distribution Q) and annealing steps (that can decrease the value of β).

Notice that, if the annealing is slow, then an ϵ -local learning algorithm with annealing can be regarded as a discrete gradient descent with respect to the joint variable (Q, β) . A natural question is does an ϵ -local learning algorithm always recover a global minimum? The following result gives a sufficient condition for an ϵ -local learning algorithm with annealing to recover the global minimum of $C_{\bar{\beta}}$.

PROPOSITION 6.4 Fix an annealing schedule $\beta_0 \geq \beta_1 \geq \dots \geq \beta_n = \bar{\beta}$. Suppose that, for every global minimizer Q_i of C_{β_i} , there exists a global minimizer Q_{i+1} of $C_{\beta_{i+1}}$ with $d(Q_i, Q_{i+1}) \leq \epsilon$. Then, an ϵ -local learning algorithm with annealing that starts from a global minimizer Q_0 of C_{β_0} , and performs one ϵ -local step after each annealing step, computes a global minimizer of $C_{\bar{\beta}}$.

We say that a task that satisfies the conditions of Theorem 6.4 for some annealing schedule is ϵ -connected. However, in general, we cannot guarantee that an ϵ -local learning algorithm cannot get stuck in a local minimum of $C_{\bar{\beta}}$. For example, this is bound to happen if there is no sequence $(Q_0, \beta_0), (Q_1, \beta_1), \dots, (Q_n, \beta_n) = (\bar{Q}, \bar{\beta})$ such that \bar{Q} is a global minimum of $C_{\bar{\beta}}$, $d(Q_i, Q_{i+1}) \leq \epsilon$, and $C_{\beta_i}(\mathcal{D}; P, Q_i) \leq C_{\beta_{i+1}}(\mathcal{D}; P, Q_{i+1})$ for all i . In particular, this happens if there is no continuous path of global minima, which is an intrinsic property of the task \mathcal{D} and loss function with respect to the function class.

6.2 SGD as a local learning algorithm

So far, we have introduced an abstract notion of distance between tasks. However, we have not yet shown that this notion is useful for DNNs, or that indeed SGD is a local algorithm in the sense of Theorem 6.2.

In [2], a step in this direction is taken. It is shown that, to first approximation, the probability of SGD converging to a configuration w_f solving task \mathcal{D} in a given time $t_f - t_0$, starting from a configuration w_0 , is given by

$$p(w_f, t_f | w_0, t_0) \simeq \underbrace{e^{-\frac{1}{2T} \Delta(L_{\mathcal{D}}(w) + \beta \text{KL}(Q \| P))}}_{\text{static part}} \underbrace{\int_{w_0}^{w_f} e^{-\frac{1}{2D} \int_{t_0}^{t_f} \frac{1}{2} \dot{u}(t)^2 + L_{\mathcal{D}}(u(t)) dt}}_{\text{dynamic part}} du(t), \quad (6.1)$$

when using the pre-distribution $P(w) \sim N(0, \lambda^2 T)$, the optimal Gaussian post-distribution $Q(w | \mathcal{D}) \sim N(w^*, \Sigma^*)$ and $\beta = 2\lambda^2 \gamma T$. Here $\Delta f(w) := f(w_f) - f(w_0)$, γ is the weight decay coefficient used to train the network, and $T \propto \eta/B$ is a temperature parameter that depends on the learning rate η and the batch size B .

From eq. 6.1 we see that, with high probability, SGD takes steps that minimize the effective potential $C_{\beta}(\mathcal{D}; P, Q) = L_{\mathcal{D}}(w) + \beta \text{KL}(Q \| P)$ (static part), while trying to minimize the distance traveled in the given time (dynamic part). Hence, SGD can be seen as a stochastic implementation of a local learning algorithm in the sense of Theorem 6.2.

In particular, this has the non-intuitive implication that SGD does not greedily optimize the loss function as it may seem from its update equation. Rather, on average, the complexity of the recovered

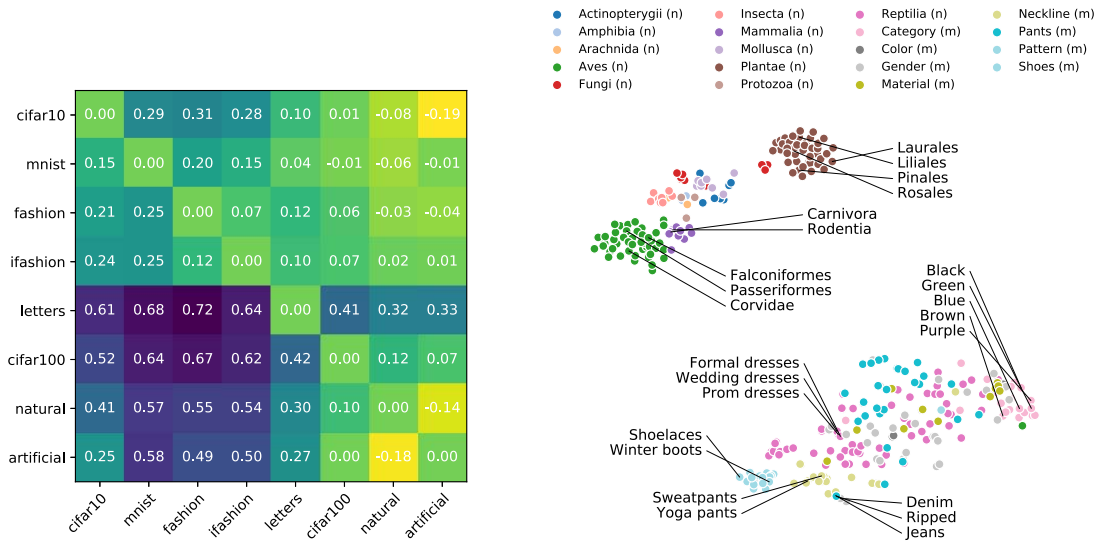


FIG. 1. **(Left)** Estimated distance matrix between several tasks. Each entry shows the distance $d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_2)$ going from the column task \mathcal{D}_1 to the row task \mathcal{D}_2 . Going from a complex task like CIFAR-100 to a simpler task (like MNIST) is always easier than the converse. Subtasks are close to the full tasks (e.g., the subset of ‘artificial’ and ‘natural’ objects of CIFAR-100 are both close to CIFAR-100). Similar tasks on the domain of small black and white images (Fashion MNIST, MNIST, Letters) are also closer together than to natural images. Inverting the colors on Fashion images leads to a very similar task (I-Fashion), as expected. **(Right)** T-SNE embedding of several organism species classification tasks and clothing attributes classification tasks based on their distance, reproduced from [1], which uses a similar definition of task distance based on the Fisher information. Intuitively, similar tasks cluster together. In the case of species classification, this largely follows the taxonomical structure.

solution affects the dynamics of the optimization process, and changes the objective.⁴ Therefore, the complexity we have introduced in Section 5 is not simply an abstract means to study different trade-offs but rather plays a concrete role in the learning dynamics.

Since SGD can be seen as a local learning algorithm, annealing the learning rate during training (i.e., annealing the parameter $\beta \propto T$) can be interpreted as a way of learning the structure of a task by slowly sweeping the structure function. Hence, SGD with annealing adaptively changes the complexity of the model, even if its dimensionality is fixed at the outset. This creates non-trivial dynamics that turns out to be beneficial to avoid memorization. It also points to the importance of the initial transient of learning during the annealing, a prime area for investigation beyond the asymptotics of the particular value of the weights, or the local structure of the residual around them, at convergence. Another consequence is that, when the task is not ϵ -connected, SGD may fail to recover the optimal structure of the data.

One may wonder if there are examples of simple tasks that are severely non- ϵ -connected, and if SGD actually fails to solve them. Biology provides such an example. [3] shows that there exist tasks—a dataset of images, and the same dataset with the same images slightly blurred—which are close to each

⁴ This was also observed by [5] although derived for the continuous approximation of SGD as a stochastic differential equation.

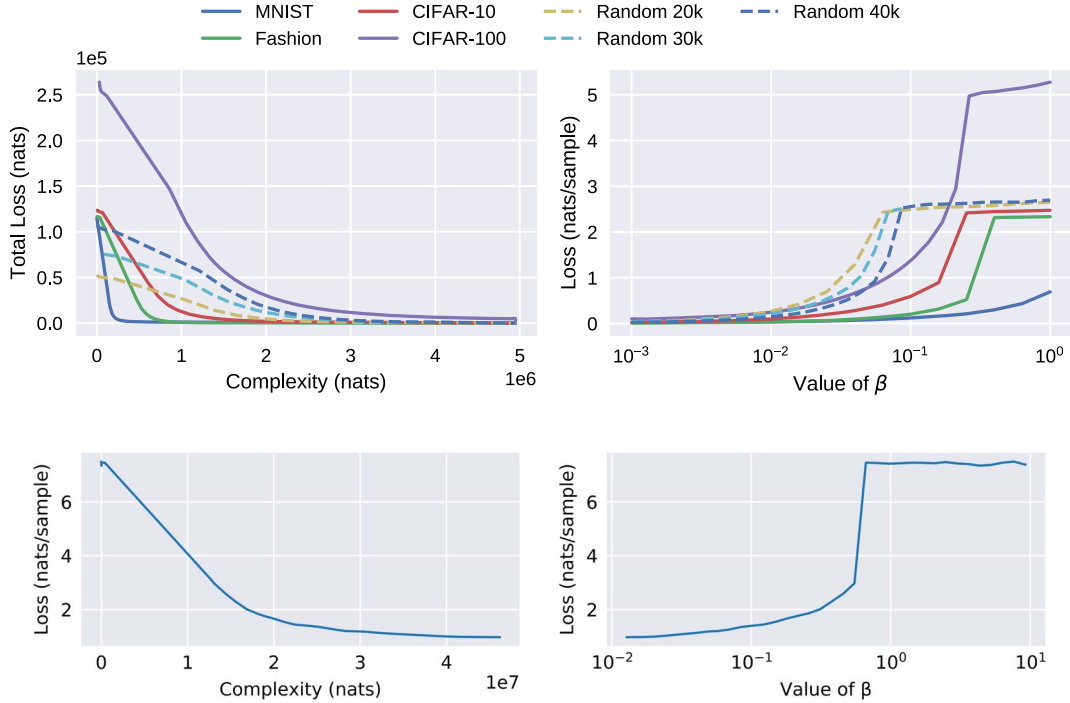


FIG. 2. We train an AllCNN model [15] by minimizing the Lagrangian $C_\beta(\mathcal{D}; P, Q)$ of eq. 5.1, for different values of β . We carry out this experiment on four real datasets (MNIST, Fashion-MNIST, CIFAR-10, CIFAR-100) and three artificial datasets of different size (20k, 30k, 40k samples) where labels are assigned randomly. In all experiments, we first pretrain for 180 epochs with weight decay $5 \cdot 10^{-4}$, batch size 128 and learning rate 0.05 (reduced by a factor of 10 after 80 and 120 epochs), by minimizing the cross-entropy loss $L_{\mathcal{D}}$. Then we fine-tune for 5 epochs, with a learning rate of 0.001 for the batch normalization parameters and for the parameters of the final classifier, and of 0.1 for all other parameters. In each epoch of fine tuning, only 10k random samples from the training set are used. **(Top left)** Approximation of the structure function $S_{\mathcal{D}}(t)$, obtained by plotting the total loss $C_\beta(\mathcal{D}; P, Q)$ against the information in the model parameters. Notice that the loss on simpler datasets rapidly decreases as we increase the complexity of the model. On the other hand, for random labels, the decrease in error is much slower and follows an almost linear trend, as described in Theorem 3.5. **(Top right)** Plot of the cross-entropy loss $L_{\mathcal{D}}$ as a function of β (which controls the trade-off between complexity and error). All datasets show a sharp transition between overfitting and underfitting at some value of β . Simpler datasets can still be fitted for a high β , since models of low complexity can already correctly classify the data. On the other hand, more complex datasets need a very low β to fit the data: In particular, random labels have the worst trade-off. The transition happens at a value that depends on the complexity of the data distribution and is mostly independent of the dataset size: the random labeled datasets all transition at a similar point, despite the difference in size. **(Bottom)** Same experiment as before, but on the ImageNet dataset [7] and using a pretrained ResNet-34 [9].

other in any intuitive sense (the taxonomical distance is zero), yet a network trained on one cannot be fine tuned to solve the other. This phenomenon, observed in diverse tasks and networks, across biological and artificial systems, is shown to be related to changes undergone by the Fisher information matrix during training. Within our framework, this can be interpreted as biasing the initial optimization process toward a minimum of $C_\beta(\mathcal{D}, P, Q)$ for the blurred task \mathcal{D} , which is not close to any of the global minima, despite being a local minimum of the Lagrangian $C_\beta(\mathcal{D}, P, Q)$. Hence, the local learning algorithm, rather than learning the structure of the data, performs sub-optimal greedy choices to optimize the loss.

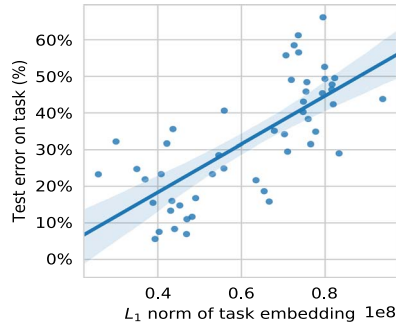


FIG. 3. **Correlation between test error and the trace of the Fisher information matrix on several tasks** (reproduced from [1]). The Fisher information matrix emerges as a complexity measure when using an uninformative prior (Section 5.1). This plot shows that the FIM trace correlates with the error obtained on the task, and so the FIM is indeed a sensible measure of the complexity of a task.

Even though the static term of eq. 6.1 is high (the distance is small), in this example the dynamic term is not negligible. This suggests the need of additional investigations into the dynamics of differential learning.

7. Empirical validation

The theoretical framework developed in this paper has tangible ramifications. A robust version of the asymmetric distance between tasks described here has been used in [1], to create a metric embedding of hundreds of real-world datasets (Fig. 1, right). The structure of the embedding shows good accordance with the complexity of the tasks (Fig. 2), and both with intuitive notions of similarity and with other metrics, such as taxonomical distance, when those are available (Fig. 1). The metric embedding allows tackling important meta-tasks, such as predicting the performance of a pre-trained model on a new datasets, given its performance on past similar datasets. This greatly simplifies the process of recommending a pre-trained model as the starting point for fine tuning a new task, which otherwise would require running a large experiment to determine the best model. The results are also shown to improve compared to pre-training a model on the most complex available task, i.e., full ImageNet classification.

Besides these independent validations, we also report some additional experiments that illustrate the concepts introduced in this paper. In Fig. 2 (right), we plot the loss as a function of β . As predicted by Fig. 3.6, the Lagrangian $C_\beta(\mathcal{D}; P, Q)$ exhibits sharp ‘phase transitions’ at a critical value β^* , when the model transitions from fitting the dataset with a trivial uniform distribution to actually fitting the labels. Regardless of their size, datasets of random labels transition around the same value β^* . The other datasets transition at a higher value, which depends on their complexity. For example, a complex dataset such a CIFAR-100 transitions at a much lower β^* than a simple dataset such as MNIST. Notice that, in this experiment, the critical value for random labels is not $\beta^* = 1$: This is because the complexity is computed using an uninformative prior (see Section 5.1), and not the universal prior as in Fig. 3.6. In Fig. 2 (left), as the complexity increases, we see that the total loss for different datasets drops to zero with different rates. This reflects the fact that MNIST, Fashion-MNIST, CIFAR-10 and CIFAR-100 are in increasing order of difficulty. Figure 2 (bottom) illustrates the results of the same experiment on the (much larger) ImageNet dataset [7].

8. Discussion and open problems

The modern practice of deep learning is rich with success stories where high-dimensional parametric models are trained on large datasets, and then adapted (fine-tuned) to specific tasks. This process requires considerable effort and expertise, and few tools are available to predict whether a given pre-trained model will succeed when fine tuned for a different task. We have started developing a language to reason about transfer learning in the abstract and analytical tools that allow to predict the success of transfer learning.

The first step is to define the tasks and the space they live in. We take a minimalistic approach and identify a task with a finite dataset. The second step is to endow the space of tasks with a metric. This is non-trivial, since different datasets can have a different cardinality and dimension, and one needs to capture the fact that a simple task is usually ‘closer’ to a complex one than vice-versa (in the sense that it is easier to fine tune a model from a complex task to a simpler one). Thus, a notion of complexity of a learning task needs to be defined first.

We introduce a notion of complexity and a notion of distance for learning tasks that are very general and encompass well-known analogues in Kolmogorov complexity theory and information theory as special cases. They espouse characteristics of Kolmogorov’s framework (focusing on finite datasets rather than asymptotics) with the intuitive notions and ease of computation of information theory. We use DNNs to compute information quantities. On one hand, this provides a convenient way of instantiating our general theory. On the other hand, it allows to measure the complexity of a deep network, and to reason about generalization (learning vs. memorization) using a non-asymptotic language, in terms of quantities that are measurable from finite datasets.

Our theory exposes connections between deep learning, complexity theory and information theory and PAC–Bayes theory, but more work is needed to fully develop these connections: the (static) distance we introduce only gives a lower bound to the feasibility of transfer learning with DNNs, and *dynamics* plays an important role as well. In particular, there are tasks that are very close, yet there is no likely path between them, so fine tuning can fail. This is observed broadly across different architectures and optimization schemes but also across different species in biology, pointing to fundamental complexity and information phenomena yet to be fully unraveled.

Acknowledgements

Work supported by ONR N00014-19-1-2229 and ARO W911NF-17-1-0304

REFERENCES

1. ACHILLE, A., LAM, M., TEWARI, R., RAVICHANDRAN, A., MAJI, S., FOWLKES, C. C., SOATTO, S. & PERONA, P. (2019) Task2Vec: task embedding for meta-learning. *Proc. IEEE Int. Conf. Comput. Vis.*, 6430–6439.
2. ACHILLE, A., MBENG, G. & SOATTO, S. (2018) The dynamics and reachability of learning tasks. arXiv preprint arXiv:1810.02440.
3. ACHILLE, A., ROVERE, M. & SOATTO, S. (2019) Critical learning periods in deep neural networks. *Proceedings of the International Conference on Learning Representations*.
4. ACHILLE, A. & SOATTO, S. (2018) Emergence of invariance and disentanglement in deep representations. *J. Mach. Learn. Res.*, **19**, 1–34.
5. CHAUDHARI, P. & SOATTO, S. (2018) Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. *2018 Information Theory and Applications Workshop (ITA)*. IEEE, pp. 1–10.
6. CSURKA, G. (2017) Domain adaptation for visual applications: a comprehensive survey. arXiv preprint arXiv:1702.05374.

7. DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K. & FEI-FEI, L. (2009) Imagenet: a large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 248–255.
8. DZIUGAITE, G. K. & ROY, D. M. (2017) Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence*.
9. HE, K., ZHANG, X., REN, S. & SUN, J. (2016) Deep residual learning for image recognition. *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 770–778.
10. HINTON, G. & VAN CAMP, D. (1993) Keeping neural networks simple by minimizing the description length of the weights. *Proceedings of the 6th Annual ACM Conference on Computational Learning Theory*.
11. LI, L. (2006) Data complexity in machine learning and novel classification algorithms. *Ph.D. Thesis*, California Institute of Technology.
12. LYGEROS, J., TOMLIN, C. & SASTRY, S. (1999) Controllers for reachability specifications for hybrid systems. *Automatica*, **35**, 349–370.
13. MARTENS, J. (2014) New insights and perspectives on the natural gradient method. arXiv preprint arXiv:1412.1193.
14. MCALLESTER, D. (2013) A PAC–Bayesian tutorial with a dropout bound. arXiv preprint arXiv:1307.2118.
15. SPRINGENBERG, J. T., DOSOVITSKIY, A., BROX, T. & RIEDMILLER, M. (2014) Striving for simplicity: the all convolutional net. arXiv preprint arXiv:1412.6806.
16. TISHBY, N., PEREIRA, F. C. & BIALEK, W. (2000) The information bottleneck method. arXiv preprint physics/0004057.
17. VERESHCHAGIN, N. K. & VITÁNYI, P. M. (2004) Kolmogorov’s structure functions and model selection. *IEEE Trans. Inf. Theory*, **50**, 3265–3290.
18. ZAMIR, A. R., SAX, A., SHEN, W., GUIBAS, L. J., MALIK, J. & SAVARESE, S. (2018) Taskonomy: disentangling task transfer learning. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3712–3722.
19. ZHANG, C., BENGIO, S., HARDT, M., RECHT, B. & VINYALS, O. (2017) Understanding deep learning requires rethinking generalization. *Proceedings of the International Conference on Learning Representations*.

A. Proofs

Proof of Theorem 3.2. (1) Let p be such that $C_K(\mathcal{D}) = L_{\mathcal{D}}(p) + K(p)$. We can compress y using $p(y|x)$, for example with an algebraic code of length $L_{\mathcal{D}}(p) = -\log p(y|x)$. A program that outputs y given x then only needs to encode the distribution $p(y|x)$ and the code for y , requiring $L_{\mathcal{D}}(p) + K(p) + O(1)$ nats, so $K(y|x) \leq C(\mathcal{D})$. For the opposite inequality, let h be the program that witness $K(y|x)$ and let $p_h(y|x) = \delta_{h(x),y}$. Then $K(p_h) = |h| = K(y|x)$ and $C_K(\mathcal{D}) \leq -\log p_h(y|x) + K(p_h) = K(y|x)$.

(2) Clearly, $C_K(\mathcal{D}) \leq C(\mathcal{D})$ as $C(\mathcal{D})$ minimizes over a smaller subset of distributions. Since $C(\mathcal{D})$ is permutation invariant, we have $C(\pi(\mathcal{D})) = C(\mathcal{D})$. Hence, $C_K(\pi(\mathcal{D})) \leq C(\pi(\mathcal{D})) = C(\mathcal{D})$.

(3) Fix a function f such that $K(f) \geq C$, and let h be a program for f . Now, consider the dataset $\mathcal{D} = (x_i, f(x_i))_{i=1}^N$, where

$$x_i = \begin{cases} \langle 0, i \rangle & \text{if } i < N \\ \langle 1, h \rangle & \text{if } i = N. \end{cases}$$

Here, the first bit is added in order to recognize the ‘special’ data point. We have $C_K(\pi(\mathcal{D})) = O(1)$ for any permutation π , because the concatenation of the input data of \mathcal{D} contains an encoding of h . On the other hand, $C(\mathcal{D}) \geq C(\mathcal{D}') \geq C$, where $\mathcal{D}' = (x_i, f(x_i))_{i=1}^{N-1}$ is the dataset obtained from \mathcal{D} by removing the special data point, and $C(\mathcal{D}') = K(f) \geq C$ if N is sufficiently large.

(4) Let $f: \mathcal{X} \rightarrow \mathcal{Y}$ be a function such that $f(x_i) = y_i$ for every $(x_i, y_i) \in \mathcal{D}$. Consider the probability distribution $p(y|x)$ defined by $p(f(x)|x) = 1$ for every $x \in \mathcal{X}$. Then we have that $K(p) \leq K(f) + O(1)$, and $L_{\mathcal{D}}(p) = 0$. Choosing f such that $K(f) = C_{\text{DET}}(\mathcal{D})$, we get $C(\mathcal{D}) \leq L_{\mathcal{D}}(p) + K(p) \leq K(f) + O(1) = C_{\text{DET}}(\mathcal{D}) + O(1)$.

To prove the equality, let $h: \{x_1, \dots, x_N\} \rightarrow N$ be the bijective function provided by the oracle. Now, create a list A of codes so that $A[h(x_i)]$ contains the encoding of y_i constructed using the distribution $p(y|x_i)$. The length of the prefix code is $\lceil -\log p(y|x_i) \rceil + 1$ (we need a prefix code so that we can concatenate all the codes). Now, given the distribution p , we can construct a function f such that $y_i = f(x_i)$ as follows: given x , compute $h(x)$, read the code $A[h(x_i)]$ and decode it using the distribution $p(y|x_i)$ to obtain the correct y_i . \square

LEMMA A1 Fix a probability distribution $p(x, y)$ on $\mathcal{X} \times \mathcal{Y}$, and assume that $p(y|x)$ is computable. Suppose that \mathcal{D} is a collection of N i.i.d. samples $(x_i, y_i) \sim p(x, y)$. Let $A_x = \{(x_i, y_i) \in \mathcal{D} \mid x_i = x\}$, and let $\hat{p}(y|x) = \frac{1}{|A_x|} \sum_{(x_i, y_i) \in A_x} \delta_{y_i, y}$ be the maximum likelihood estimation (MLE) of $p(y|x)$. Then, for every $\epsilon > 0$, there exists $c > 0$ such that, in the limit $N \rightarrow \infty$ we have

$$|L_{\mathcal{D}}(p) - L_{\mathcal{D}}(\hat{p})| < c$$

with probability $> 1 - \epsilon$.

Proof of Lemma A1. Consider $p(y|x)$ and $\hat{p}(y|x)$ as vectors of size $|\mathcal{X} \times \mathcal{Y}|$. Expand $L_{\mathcal{D}}(p)$ at \hat{p} :

$$L_{\mathcal{D}}(p) - L_{\mathcal{D}}(\hat{p}) = \nabla_p L_{\mathcal{D}}(\hat{p})(p - \hat{p}) + (p^* - \hat{p})^T \nabla_p^2 L_{\mathcal{D}}(\hat{p})(p^* - \hat{p}).$$

Where p^* is on the line connecting p and \hat{p} . The first term is zero, since the MLE estimation \hat{p} is by definition a minimum of $L_{\mathcal{D}}$. Now, recall that the MLE converges to the real distribution as $\sqrt{N}(\hat{p} - p) \rightarrow N(0, F(p)^{-1})$, where $F(p)$ is the Fisher information matrix computed in p , and $\nabla_p^2 L_{\mathcal{D}}(\hat{p}) < N(\nabla_p^2 L_{\mathcal{D}}(p) + c')$. Let $a > 0$ be such that, with probability $1 - \epsilon$, $\|\hat{p} - p\| < a$. Then,

$$|L_{\mathcal{D}}(p) - L_{\mathcal{D}}(\hat{p})| = |(p^* - \hat{p})^T \nabla_p^2 L_{\mathcal{D}}(\hat{p})(p^* - \hat{p})| < a^2 \text{tr}(\nabla_p^2 L_{\mathcal{D}}(p) + c') = c.$$

\square

Proof of Theorem 3.3. (i) By Shannon's coding theorem, the expected value of $K(y|x)$ is at least $N \cdot H_p(y|x)$. Then, the first inequality follows from part (1) of Theorem 3.2. If we use $p(y|x)$ itself in the definition of $C(\mathcal{D})$, we obtain $C(\mathcal{D}) \leq L_{\mathcal{D}}(p) + K(p(y|x))$. The expected value of $L_{\mathcal{D}}(p)$ is $N \cdot H_p(y|x)$, so we obtain the second inequality.

(ii) We need to prove that, for any other distribution p' , we have

$$L_{\mathcal{D}}(p') + K(p') > L_{\mathcal{D}}(p) + K(p).$$

In fact, suppose that $L_{\mathcal{D}}(p') + K(p') \leq L_{\mathcal{D}}(p) + K(p)$. Then we have

$$L_{\mathcal{D}}(p') - L_{\mathcal{D}}(p) \leq K(p) - K(p').$$

Notice that we can lower bound the LHS using the MLE estimator \hat{p} , which by definition minimizes $L_{\mathcal{D}}$. By Lemma A1, we have $L_{\mathcal{D}}(\hat{p}) - L_{\mathcal{D}}(p) \geq -c$, hence $K(p') \leq K(p) + c$. By the central limit theorem, the LHS grows as $N \cdot \mathbb{E}_x[\text{KL}(p(y|x) \| p'(y|x))] + O(\sqrt{N})$, so we must have

$$\mathbb{E}_x[\text{KL}(p(y|x) \| p'(y|x))] \leq \frac{K(p)}{N} + O\left(\frac{1}{\sqrt{N}}\right).$$

Therefore, $\mathbb{E}_x[\text{KL}(p(y|x) \| p'(y|x))] \rightarrow 0$ as $N \rightarrow \infty$. On the other hand, we have that

$$\mathbb{E}_x[\text{KL}(p(y|x) \| p'(y|x))] \geq k > 0$$

for some k , since there are only finitely many distributions $p' \neq p$ such that $K(p') \leq K(p) + c$. \square

Proof of Lemma 4.1. Positivity follows trivially by the definition. The second property is also easy: $d_{\beta}(\mathcal{D} \rightarrow \mathcal{D}) \leq \max_p K(p|p) = O(1)$.

We now prove the triangle inequality. In what follows, p_i and \bar{p}_i are always β -minimal sufficient statistics of \mathcal{D}_i . Fix any \bar{p}_1 . Choose \bar{p}_2 such that $K(\bar{p}_2|\bar{p}_1)$ is minimized. Then choose \bar{p}_3 such that $K(\bar{p}_3|\bar{p}_2)$ is minimized. Then

$$\begin{aligned} \min_{p_3} K(p_3|\bar{p}_1) &\leq K(\bar{p}_3|\bar{p}_1) \\ &\leq K(\bar{p}_2|\bar{p}_1) + K(\bar{p}_3|\bar{p}_2) + O(1) \\ &\leq d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_2) + d_{\beta}(\mathcal{D}_2 \rightarrow \mathcal{D}_3) + O(1). \end{aligned}$$

This holds for every \bar{p}_1 . Taking the maximum over \bar{p}_1 , we obtain the desired result. \square

Proof of Theorem 4.3. For $(x, y) \in \mathcal{X} \times \mathcal{Y}$ and $i \in \{1, 2\}$, define $p_{12}(y|x, i) = p_i(y|x)$. Then $\mathcal{D}_1 \sqcup \mathcal{D}_2$ is effectively obtained by sampling from p_{12} . By Theorem 3.3, with high probability and for $|\mathcal{D}_1|$ and $|\mathcal{D}_2|$ sufficiently large, we have that p_1 and p_{12} are the unique minimal sufficient statistics for \mathcal{D}_1 and $\mathcal{D}_1 \sqcup \mathcal{D}_2$, respectively. By definition of p_{12} , we have that $K(p_1|p_{12}) = O(1)$. Then $d_{\beta}(\mathcal{D}_1 \sqcup \mathcal{D}_2 \rightarrow \mathcal{D}_1) = O(1)$. \square

Proof of Corollary 4.1. One inequality follows by Lemma 4.1 and Theorem 4.3:

$$\begin{aligned} d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_2) &\leq d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_1 \sqcup \mathcal{D}_2) + d_{\beta}(\mathcal{D}_1 \sqcup \mathcal{D}_2 \rightarrow \mathcal{D}_2) + O(1) \\ &= d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_1 \sqcup \mathcal{D}_2) + O(1). \end{aligned}$$

To prove the other inequality, define p_{12} as in the proof of Theorem 4.3. Then

$$\begin{aligned} d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_1 \sqcup \mathcal{D}_2) &= K(p_{12}|p_1) \\ &\leq K(p_2|p_1) + O(1) \\ &= d_{\beta}(\mathcal{D}_1 \rightarrow \mathcal{D}_2) + O(1). \end{aligned}$$

\square

Proof of Theorem 5.2. We have $\text{KL}(Q(w|\mathcal{D}) \| P(w)) = -\log(e^{-K(w^*)}/Z) = K(w^*) + \log(Z)$. \square

Proof of Theorem 5.3. For a fixed training algorithm $A : \mathcal{D} \mapsto Q(w|\mathcal{D})$, we want to find the prior $P^*(w)$ that minimizes the expected complexity of the data:

$$\begin{aligned} P^*(w) &= \operatorname{argmin}_{P(w)} \mathbb{E}_{\mathcal{D}}[C(\mathcal{D})] \\ &= \operatorname{argmin}_{P(w)} \left[\mathbb{E}_{\mathcal{D}}[L_{\mathcal{D}}(p_w(y|x))] + \mathbb{E}_{\mathcal{D}}[\operatorname{KL}(Q(w|\mathcal{D}) \| P(w))] \right]. \end{aligned}$$

Notice that only the second term depends on $P(w)$. Let $Q(w) = \mathbb{E}_{\mathcal{D}}[Q(w|\mathcal{D})]$ be the marginal distribution of w , averaged over all possible training datasets. We have

$$\mathbb{E}_{\mathcal{D}}[\operatorname{KL}(Q(w|\mathcal{D}) \| P(w))] = \mathbb{E}_{\mathcal{D}}[\operatorname{KL}(Q(w|\mathcal{D}) \| Q(w))] + \mathbb{E}_{\mathcal{D}}[\operatorname{KL}(Q(w) \| P(w))].$$

Since the KL divergence is always positive, the optimal ‘adapted’ prior is given by $P^*(w) = Q(w)$, i.e., the marginal distribution of w over all datasets. Finally, by definition of Shannon’s mutual information, we get

$$I(w; \mathcal{D}) = \operatorname{KL}(Q(w|\mathcal{D}) \pi(\mathcal{D}) \| Q(w) \pi(\mathcal{D})) = \mathbb{E}_{\mathcal{D} \sim \pi(\mathcal{D})}[\operatorname{KL}(Q(w|\mathcal{D}) \| Q(w))]. \quad \square$$

Proof of Theorem 5.4. Since both $P(w)$ and $Q(w|\mathcal{D})$ are Gaussian distributions, the KL divergence can be written as

$$\operatorname{KL}(Q(w|\mathcal{D}) \| P(w)) = \frac{1}{2} \left[\frac{\|\mu\|^2}{\lambda^2} + \frac{1}{\lambda^2} \operatorname{tr}(\Sigma) + k \log \lambda^2 - \log |\Sigma| - k \right],$$

where k is the number of components of w .

Let w^* be a local minimum of the cross-entropy loss $L_{\mathcal{D}}(p_w(y|x))$, and let H be the Hessian of $L_{\mathcal{D}}(p_w(y|x))$ in w^* . Set $\mu = w^*$. Assuming that a quadratic approximation holds in a sufficiently large neighborhood, we obtain

$$C_{\beta}(\mathcal{D}; P, Q) = L_{\mathcal{D}}(p_{w^*}(y|x)) + \frac{1}{2} \operatorname{tr}(H \cdot \Sigma) + \frac{\beta}{2} \left[\frac{\|w^*\|^2}{\lambda^2} + \frac{1}{\lambda^2} \operatorname{tr}(\Sigma) + k \log \lambda^2 - \log |\Sigma| - k \right].$$

The gradient with respect to Σ is

$$\frac{\partial C_{\beta}(\mathcal{D}; P, Q)}{\partial \Sigma} = \frac{1}{2} \left[H + \frac{\beta}{\lambda^2} I - \beta \Sigma^{-1} \right]^{\top}.$$

Setting it to zero, we obtain the minimizer $\Sigma^* = \beta(H + \frac{\beta}{\lambda^2} I)^{-1}$.

Recall that the Hessian of the cross-entropy loss coincides with the Fisher information matrix F at w^* , because w^* is a critical point [13]. Since $L_{\mathcal{D}}(p_w(y|x))$, and hence H , is not normalized by the number of samples N , the exact relation is $H = N \cdot F$. Taking the limit for $\lambda \rightarrow \infty$, we obtain the desired result. \square

Proof of Theorem 6.4. If Q_i is a global minimizer of C_{β_i} , then it is at distance at most ϵ from a global minimizer Q_{i+1} of $C_{\beta_{i+1}}$. Therefore, for $\beta = \beta_{i+1}$, an ϵ -local step from Q_i reaches a global minimizer Q_{i+1} of $C_{\beta_{i+1}}$. By induction, the algorithm terminates at a global minimizer \bar{Q} of $C_{\bar{\beta}}$. \square