
Frustratingly Easy Transferability Estimation

Long-Kai Huang¹ Junzhou Huang¹ Yu Rong¹ Qiang Yang² Ying Wei³

Abstract

Transferability estimation has been an essential tool in selecting a pre-trained model and the layers in it for transfer learning, so as to maximize the performance on a target task and prevent negative transfer. Existing estimation algorithms either require intensive training on target tasks or have difficulties in evaluating the transferability between layers. To this end, we propose a simple, efficient, and effective transferability measure named TransRate. Through a single pass over examples of a target task, TransRate measures the transferability as the mutual information between features of target examples extracted by a pre-trained model and their labels. We overcome the challenge of efficient mutual information estimation by resorting to coding rate that serves as an effective alternative to entropy. From the perspective of feature representation, the resulting TransRate evaluates both completeness (whether features contain sufficient information of a target task) and compactness (whether features of each class are compact enough for good generalization) of pre-trained features. Theoretically, we have analyzed the close connection of TransRate to the performance after transfer learning. Despite its extraordinary simplicity in 10 lines of codes, TransRate performs remarkably well in extensive evaluations on 32 pre-trained models and 16 downstream tasks.

1. Introduction

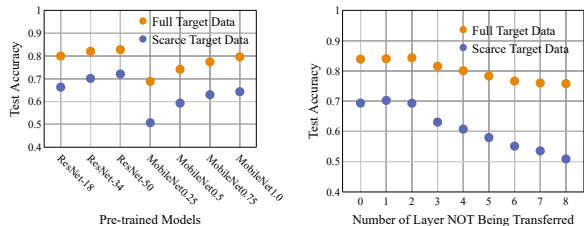
Transfer learning from standard large datasets (e.g., ImageNet) and corresponding pre-trained models (e.g., ResNet-50) has become a de-facto method for real-world deep learning applications where limited annotated data is accessible. Unfortunately, the performance gain by transfer learning

could vary a lot, even with the possibility of negative transfer (Pan & Yang, 2009; Wang et al., 2019; Zhang et al., 2020). First, the *relatedness of the source task* where a pre-trained model is trained on to the target task largely dictates the performance gain. Second, using pre-trained models in different *architectures* also leads to uneven performance gain, even for the same pair of source and target tasks. Figure 1(a) tells that ResNet-50 pre-trained on ImageNet contributes the most to the target task CIFAR-100, compared to the other architectures. Finally, the optimal *layers* to transfer vary from pair to pair. While higher layers encode more semantic patterns that are specific to source tasks, lower-layer features are more generic (Yosinski et al., 2014). Especially if a pair of tasks are not sufficiently similar, determining the optimal layers is expected to strike the balance between transferring only lower-layer features (as higher layers specific to a source task may hurt the performance of a target task) and transferring more higher-layer features (as training more higher-layers from scratch requires extensive labeled data). As shown in Figure 1(b), not transferring the three highest layers is preferred for training with full target data, though transferring all but the two highest layers achieves the highest test accuracy with scarce target data. This suggests the following

Question: Which pre-trained model (possibly trained on different source tasks in a supervised or unsupervised manner) and which layers of it should be transferred to benefit the target task the most?

This research question drives the design of transferability estimation methods, including computation-intensive (Achille et al., 2019; Dwivedi & Roig, 2019; Song et al., 2020; Zamir et al., 2018) and computation-efficient ones (Bao et al., 2019; Cui et al., 2018; Nguyen et al., 2020; Tran et al., 2019; You et al., 2021). The pioneering works (Achille et al., 2019; Zamir et al., 2018) directly follow the definition of transfer learning to measure the transferability, and thereby require fine-tuning on a target task with expensive parameter optimization. Though their follow-ups (Dwivedi & Roig, 2019; Song et al., 2020) alleviate the need of fine-tuning, their prerequisites still include an encoder pre-trained on target tasks. Keeping in mind that the primary goal of a transferability measure is to select a pre-trained model prior to training on a target task, researchers recently turned towards computation-efficient ways. The transferability is estimated

¹Tencent AI Lab ²Hong Kong University of Science and Technology ³City University of Hong Kong. Correspondence to: Ying Wei <yingwei@cityu.edu.hk>.



(a) Test accuracy for pre-trained ImageNet with different layers of the pre-trained model architectures. (b) Test accuracy for transfer-trained ResNet34 model.

Figure 1: Transferring from ImageNet to CIFAR-100. For “Full Target Data”, all target data are used, while for “Scarce Target Data”, only 50 target samples per class are used in training.

as the negative conditional entropy between the labels of the two tasks in (Tran et al., 2019; Nguyen et al., 2020). Bao et al. (2019) and You et al. (2021) solved two surrogate optimization problems to estimate the likelihood and marginalized likelihood of labeled target examples, under the assumption that a linear classifier is added on top of the pre-trained model. However, the efficiency comes at the price of failing to discriminate transferability between layers, for (Nguyen et al., 2020; Tran et al., 2019) that estimate with labels only and (Bao et al., 2019; You et al., 2021) that consider transferring the penultimate layer only.

We are motivated to pursue a computation-efficient transferability measure without sacrificing the merit of computation-intensive methods in comprehensive transferability evaluation, especially between layers. Mutual information between features and labels has a strong predictive role in the effectiveness of feature representation, dated back to the decision tree algorithm (Quinlan, 1986) and also evidenced in recent studies (Tishby & Zaslavsky, 2015). Markedly, mutual information varies from layer to layer of features given labels, which makes itself an attractive measure for transferability. In this paper, we propose to estimate the transferability with the mutual information between labels and features of target examples extracted by a pre-trained model at a specific layer. Though mutual information itself is notoriously challenging to estimate (Hjelm et al., 2019), we overcome this obstacle by resorting to the coding rate proposed in (Ma et al., 2007) inspired from the close connection between rate distortion and entropy in information theory (Cover, 1999). The resulting estimation named TransRate offers the following advantages: 1) it perfectly matches our need in computation efficiency, free of either prohibitively exhaustive discretization (Tishby & Zaslavsky, 2015) or neural network training (Belghazi et al., 2018); 2) it is well defined for finite examples from a subspace-like distribution, even if the examples are represented in a high dimensional feature space of a pre-trained model. In a nutshell, TransRate is simple and efficient, as the only computations it requires are (1) making a single forward

pass of the model pre-trained on a source task through the target examples to obtain their features at a set of selected layers and (2) calculating the TransRate.

Despite being “frustratingly easy”, TransRate enjoys the following benefits that we would highlight.

- TransRate allows to select between layers of a pre-trained model for better transfer learning.
- We have theoretically analyzed that TransRate closely aligns with the performance after transfer learning.
- A larger value of TransRate is associated with more complete and compact features that are strongly suggestive of better generalization.
- TransRate offers surprisingly good performance on transferability comparison between source tasks, between architectures, and between layers. We investigate a total of 32 pre-trained models (including supervised, unsupervised, self-supervised, convolutional and graph neural networks), 16 downstream tasks (including classification and regression), and the insensitivity of TransRate against the number of labeled examples in a target task.

2. Related Works

Re-training or fine-tuning a pre-trained model is a simple yet effective strategy in transfer learning (Pan & Yang, 2009). To improve the performance on a target task and avoid negative transfer, there have been various works on transferability estimation between tasks (Achille et al., 2019; Bao et al., 2019; Cui et al., 2018; Dwivedi & Roig, 2019; Nguyen et al., 2020; Song et al., 2020; Tran et al., 2019; Zamir et al., 2018; Li et al., 2021), which we summarize them in Table 1. Taskonomy (Zamir et al., 2018) and Task2Vec (Achille et al., 2019) evaluate the task relatedness by the loss and the Fisher Information Matrix after fully performing fine-tuning of the pre-trained model on the target task, respectively. In RSA (Dwivedi & Roig, 2019) and DEPARA (Song et al., 2020), the authors proposed to build a similarity graph between examples for each task based on representations by a pre-trained model on this task, and took the graph similarity across tasks as the transferability. Despite their general applicability in using unsupervised pre-trained models besides supervised ones and selecting the layer to transfer, their computational costs that are as high as fine-tuning with target labeled data exclude their applications to meet the urgent need of transferability estimation prior to fine-tuning.

There also exist transferability measures proposed for domain generalization (Zhang et al., 2021) and multi-source transfer (Tong et al., 2021), where a special class of integral probability metric between domains and the optimal combination coefficients of source models that minimizes the χ^2 between the combined source distribution and the target distribution were proposed, respectively. Both of them stand in need of source datasets; however, we focus on evaluat-

Table 1: Summary of the existing transferability measures and ours.

Measures	Free of Training on Target	Free of Assessing Source	Free of Optimization	Applicable to Unsupervised Pre-trained Models	Applicable to Layer Selection
Taskonomy (Zamir et al., 2018)	×	×	✓	✓	✓
Task2Vec (Achille et al., 2019)	×	×	×	✓	×
RSA (Dwivedi & Roig, 2019)	×	✓	✓	✓	✓
DEPARA (Song et al., 2020)	×	✓	✓	✓	✓
\mathcal{N} LEEP (Li et al., 2021)	×	✓	✓	✓	✓
DS (Cui et al., 2018)	✓	×	×	✓	×
(Zhang et al., 2021)	✓	×	×	×	×
(Tong et al., 2021)	✓	×	×	×	×
NCE (Tran et al., 2019)	✓	×	✓	×	×
H-Score (Bao et al., 2019)	✓	✓	×	✓	×
LogME (You et al., 2021)	✓	✓	×	✓	×
LEEP (Nguyen et al., 2020)	✓	✓	✓	×	×
TransRate	✓	✓	✓	✓	✓

ing the transferability of various pre-trained models, where the source dataset that a pre-trained model is trained on is oftentimes too huge and private to access.

This work is more aligned with recent attempts towards computationally efficient transferability measures without training on target data (Bao et al., 2019; Cui et al., 2018; Nguyen et al., 2020; Tran et al., 2019; You et al., 2021). The Earth Mover’s Distance between features of the source and the target is used in (Cui et al., 2018). Tran et al. (2019) proposed the NCE score to estimate the transferability by the negative conditional entropy between labels of a target and a source task. But alas, the reliance on source datasets again disable these two methods towards assessing the transferability of a broad range of pre-trained models. To bypass the limitations, Bao et al. (2019) and You et al. (2021) proposed to directly estimate the likelihood and the marginalized likelihood of labeled target examples, respectively, by assuming that a linear classifier is added on top of the pre-trained model. Nguyen et al. (2020) proposed the LEEP score, where source labels used in NCE (Tran et al., 2019) are replaced with soft labels generated by the pre-trained model. Its extension (Li et al., 2021) computes more accurate soft source labels via a fitted a Gaussian mixture model (GMM), at the undesirable cost of training the GMM on the target set similar to computation-intensive methods. Neither of the three, however, is designed for layer selection – H-Score (Bao et al., 2019) and LogME (You et al., 2021) consider the penultimate layer to be transferred only and LEEP estimating the transferability with labels only fails to differentiate by layers. Our purpose of the proposed TransRate is a simple but effective transferability measure: 1) it is optimization-free with single forward pass, without solving optimization problems as in (Bao et al., 2019; You et al., 2021); 2) besides selecting the source and the architecture of a pre-trained model, it supports layer selection to fill the gap in computationally-efficient measures.

3. TransRate

3.1. Notations and Problem Settings

We consider the knowledge transfer from a source task T_s to a target task T_t of C -category classification. As widely accepted, only the model that is pre-trained on the source task, instead of source data, is accessible. The pre-trained model, denoted by $F = f_{L+1} \circ \dots \circ (f_2 \circ f_1)$, consists of an L -layer feature extractor and a 1-layer classifier f_{L+1} . Here, f_l is the mapping function at the l -th layer. The target task is represented by n labeled data samples $\{(x_i, y_i)\}_{i=1}^n$. Afterwards, we denote the number of layers to be transferred by K ($K \leq L$). These K layers of the model are named as the pre-trained feature extractor $g = f_K \circ \dots \circ (f_2 \circ f_1)$. The feature of x_i extracted by g is denoted as $z_i = g(x_i)$. Building on the feature extractor, we construct the target model denoted by w to include 1) the same structure as the $(K+1)$ -th to (L) -th layers of the source model and 2) a new classifier f_{L+1}^t for the target task. We also refer to w as the head of the target model. Following the standard practice of fine-tuning, both the feature extractor g and the head w will be trained on the target task.

We consider the optimal model for the target task as

$$g^*, w^* = \arg \max_{\tilde{g} \in \mathcal{G}, w \in \mathcal{W}} \mathcal{L}(\tilde{g}, w) = \arg \max_{\tilde{g} \in \mathcal{G}, w \in \mathcal{W}} \frac{1}{n} \sum_{i=1}^n \log p(y_i | z_i; \tilde{g}, w)$$

subject to $\tilde{g}^{(0)} = g$, where \mathcal{L} denotes the log-likelihood, and \mathcal{G} and \mathcal{W} are the spaces of all possible feature extractors and heads, respectively. We define the transferability as the expected log-likelihood of the optimal model $w^* \circ g^*$ on test samples in the target task:

Definition 1 (Transferability). *The transferability of a pre-trained feature extractor g from a source task T_s to a target task T_t , denoted by $\text{Trf}_{T_s \rightarrow T_t}(g)$, is measured by the expected log-likelihood of the optimal model $w^* \circ g^*$*

on a random test sample (x, y) of T_t : $\text{Trf}_{T_s \rightarrow T_t}(g) := \mathbb{E}[\log p(y|z^*; g^*, w^*)]$ where $z^* = g^*(x)$.

This definition of transferability can be used for 1) selection of a pre-trained feature extractor among a model zoo $\{g_m\}_{m=1}^M$ for a target task, where M pre-trained models could be in different architectures and trained on different source tasks in a supervised or unsupervised manner, and 2) selection of a layer to transfer among all configurations $\{g_m^l\}_{l=1}^K$ given a pre-trained model g_m and a target task.

3.2. Computation-Efficient Transferability Estimation

Computing the transferability defined in Definition 1 is as prohibitively expensive as fine-tuning all M pre-trained models or K layer configurations of a pre-trained model on the target task, while the transferability offers benefits only when it can be calculated a priori. To address this shortfall, we propose TransRate, a frustratingly easy measure, to estimate the defined transferability. The transferability characterizes how well the optimal model, composed of the feature extractor g^* initialized from g and the head w^* , performs on the target task, where the performance is evaluated by the log-likelihood. However, the optimal model $w^* \circ g^*$ is inaccessible without optimizing $\mathcal{L}(\tilde{g}, w)$. For tractability, we follow prior computation-efficient transferability measures (Nguyen et al., 2020; You et al., 2021) to estimate the performance of $w^* \circ g$ instead. By reasonably assuming that w^* can extract all the information related to the target task in the pre-trained feature extractor g , we argue that the mutual information between the pre-trained feature extractor g and the target task serves as a strong indicator for the performance of the model $w^* \circ g$. Therefore, the proposed TransRate measures this mutual information as,

$$\text{TrR}_{T_s \rightarrow T_t}(g) = h(Z) - h(Z|Y) \approx H(Z^\Delta) - H(Z^\Delta|Y), \quad (1)$$

where Y are labels of target examples, $Z = g(X)$ and Z^Δ are features and quantized features of them extracted by the pre-trained feature extractor g . Eqn. (1) follows $h(Z) \approx H(Z^\Delta) + \log \Delta$ ($\Delta \rightarrow 0$) (Cover, 1999), where $H(\cdot)$ denotes the Shannon entropy of a discrete random variable (e.g., Z^Δ with the quantization error Δ), and $h(\cdot)$ is the differential entropy of a continuous random variable (e.g., Z).

Based on the theory in (Qin et al., 2019), we show in Proposition 1 that TransRate provides an upper bound and lower bound to the log-likelihood of the model $w^* \circ g$.

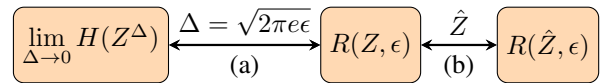
Proposition 1. *Assume the target task has a uniform label distribution, i.e., $p(Y = y^c) = \frac{1}{C}$ holds for all $c = 1, 2, \dots, C$. We then have:*

$$\begin{aligned} \mathcal{L}(g, h^*) &\lesssim \text{TrR}_{T_s \rightarrow T_t}(g) - H(Y), \\ \mathcal{L}(g, h^*) &\gtrsim \text{TrR}_{T_s \rightarrow T_t}(g) - H(Y) - H(Z^\Delta). \end{aligned}$$

Note that NCE, LEEP and TransRate all provide a lower bound for the maximal log-likelihood, whereas only TransRate has been shown to be a tight upper bound of the maximal log-likelihood. Since the maximal log-likelihood is closely related to the transfer learning performance, this proposition implies that TransRate highly aligns with the transfer performance. A detailed proof and more analysis on the relationship between TransRate and transfer performance can be found in Appendix D.1 and Appendix C.

Computing the TransRate in Eqn. (1), however, remains a daunting challenge, as the mutual information is notoriously difficult to compute especially for continuous variables in high-dimensional settings (Hjelm et al., 2019). A popular solution for mutual information estimation is to have the quantization Z^Δ via the histogram method (Tishby & Zaslavsky, 2015), though it requires an extremely large memory capacity. Even if we divide each dimension of Z into only 10 bins, there will be 10^d bins where d is the dimension of Z that is usually greater than 128. Other estimators include kernel density estimator (Moon et al., 1995) and k-NN estimator (Beirlant et al., 1997; Kraskov et al., 2004). KDE suffers from singular solutions when the number of examples is smaller than their dimension; the k -NN estimator requiring exhaustive computation of nearest neighbors of all examples may be too computationally expensive if more examples are available. Recent trends in deep neural networks have led to a proliferation studies of approximating the mutual information or entropy by a neural network (Belghazi et al., 2018; Hjelm et al., 2019; Shalev et al., 2020) and obtaining a high-accuracy estimation by optimizing the neural network. Unfortunately, training neural networks is contrary to our premise of an optimization-free transferability measure.

Fortunately, as shown in Figure 2, the rate distortion $R(Z, \epsilon)$ defining the minimal number of binary bits to encode Z with an expected decoding error less than ϵ has been proved to be closely related to the Shannon entropy, i.e., $R(Z, \epsilon) = H(Z^\Delta) + o(1)$ with $\Delta = \sqrt{2\pi e \epsilon}$ when $\epsilon \rightarrow 0$ (Binia et al., 1974; Cover, 1999). Most crucially, the work of (Ma et al.,



Shannon entropy Rate distortion Coding rate

Figure 2: Illustration of the relationship between the three information measures: (a) the rate distortion of a continuous random variable amounts to $H(Z^{\sqrt{2\pi e \epsilon}}) + o(1)$ when $\epsilon \rightarrow 0$ (Binia et al., 1974), where a larger ϵ introduces an approximation error; (b) the coding rate provides an empirical estimate of the rate distortion, where the approximation error is dictated by the degree to which finite samples \hat{Z} represent the true random variable Z .

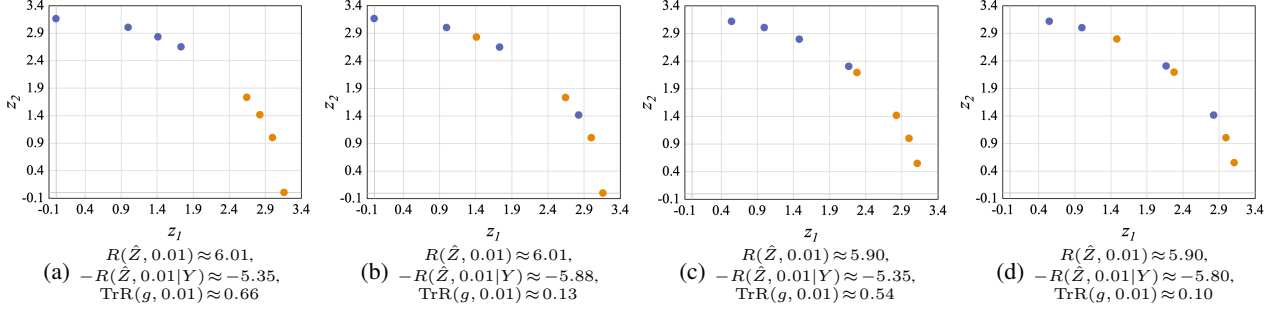


Figure 3: Toy examples illustrating the effectiveness of the TransRate. The horizontal and vertical axes represent the two dimensions of features \hat{Z} . There are two classes in Y , pictorially illustrated with two colors.

2007) offers the coding rate $R(\hat{Z}, \epsilon)$ as an efficient and accurate empirical estimation to $R(Z, \epsilon)$, provided with n finite samples $\hat{Z} = [z_1, z_2, \dots, z_n] \in \mathbb{R}^{d \times n}$ from a subspace-like distribution where d is the dimension of z_i . Concretely,

$$R(\hat{Z}, \epsilon) = \frac{1}{2} \log \det(I_d + \frac{1}{n\epsilon} \hat{Z} \hat{Z}^\top), \quad (2)$$

where ϵ is the distortion rate. Coding rate has been verified to be qualified even for samples that are in high-dimensional feature representations or from a non-Gaussian distribution (Ma et al., 2007), which is often the case for features by deep neural networks. Therefore, we resort to $R(\hat{Z}, \epsilon)$ as an approximation to $H(Z^\Delta)$ ($\Delta = \sqrt{2\pi\epsilon}$) with a small value of ϵ . More properties of the coding rate will be discussed in Appendix C.1 and Appendix D.2.

Next we investigate the rate distortion estimate $R(\hat{Z}, \epsilon|Y)$ as an approximation to the second component of TransRate, i.e., $H(Z^\Delta|Y)$. Define $Z^c = \{z|Y = y^c\}$ as the random variable for features of the target samples in the c -th class, whose labels are all y^c . When $\epsilon \rightarrow 0$, we then have

$$\begin{aligned} H(Z^\Delta|Y) &\approx h(Z|Y) - \log \Delta \\ &= - \int \sum_{z \in Z} p(z, y) \log p(z|y) dz - \log \Delta \\ &= - \int \sum_{z \in Z} \sum_{c=1}^C p(Y = y^c) p(z|Y = y^c) \log p(z|Y = y^c) dz - \log \Delta \\ &= \sum_{c=1}^C \frac{n_c}{n} \int_{z \in Z^c} -p(z) \log p(z) dz - \log \Delta \\ &= \sum_{c=1}^C \frac{n_c}{n} [h(Z^c) - \log \Delta] = \sum_{c=1}^C \frac{n_c}{n} H((Z^c)^\Delta), \end{aligned} \quad (3)$$

where n_c is the number of training samples in the c -th class. According to (3), it is direct to derive

$$\begin{aligned} R(\hat{Z}, \epsilon|Y) &= \sum_{c=1}^C \frac{n_c}{n} R(\hat{Z}^c, \epsilon) \\ &= \sum_{c=1}^C \frac{n_c}{2n} \log \det(I_d + \frac{1}{n_c \epsilon} \hat{Z}^c \hat{Z}^{c\top}), \end{aligned} \quad (4)$$

where $\hat{Z}^c = [z_1^c, z_2^c, \dots, z_{n_c}^c] \in \mathbb{R}^{d \times n_c}$ denotes n_c samples in the c -th class. Combining (2) and (4), we conclude with the TransRate we use in practice for transferability estimation:

$$\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon) = R(\hat{Z}, \epsilon) - R(\hat{Z}, \epsilon|Y).$$

Note that we use $\text{TrR}_{T_s \rightarrow T_t}(g)$ and $\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon)$ to denote the ideal and the working TransRate, respectively.

Completeness and Compactness We argue that those pre-trained models that produce both complete and compact features tend to have high TransRate scores. (1) *Completeness*: $R(\hat{Z}, \epsilon)$ as the first term of $\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon)$ evaluates whether the features \hat{Z} by the pre-trained feature extractor g include sufficient information for solving the target task – features between different classes of examples should be as diverse as possible. \hat{Z} in Figure 3(a) is more diverse than that in Figure 3(c), evidenced by a larger value of $R(\hat{Z}, 0.01)$. (2) *Compactness*: The second term, i.e., $-R(\hat{Z}, \epsilon|Y)$, assesses whether the features \hat{Z}^c for each c -th class are compact enough for good generalization. Each of the two classes spans a wider range in Figure 3(b) than that in Figure 3(a), so that the value of $-R(\hat{Z}, 0.01|Y)$ is smaller.

Furthermore, there is theoretical evidence to strengthen the argument above. Consider a binary classification problem with $\hat{Z} = [\hat{Z}^1, \hat{Z}^2] \in \mathbb{R}^{d \times n}$, where both \hat{Z}^1 and \hat{Z}^2 have n d -dim examples. By defining $\alpha = 1/n\epsilon$, we have $\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon) = \frac{1}{2} \log \det\{(I_{n/2} + \alpha(\hat{Z}^1)^\top \hat{Z}^1 + \alpha(\hat{Z}^2)^\top \hat{Z}^2) + \alpha^2[(\hat{Z}^1)^\top \hat{Z}^1 (\hat{Z}^2)^\top \hat{Z}^2 - (\hat{Z}^1)^\top \hat{Z}^2 (\hat{Z}^2)^\top \hat{Z}^1]\} - B$ where $B = \frac{1}{2}(R(\hat{Z}^1, \epsilon) + R(\hat{Z}^2, \epsilon))$. We assume $(\hat{Z}^1)^\top \hat{Z}^1$ and $(\hat{Z}^2)^\top \hat{Z}^2$ to be fixed, so that TransRate maintaining the *compactness* within each class (i.e., B is a constant) and hinging on the *completeness* only maximizes at $(\hat{Z}^1)^\top \hat{Z}^2 = 0$ and minimizes at $\hat{Z}^1 = \hat{Z}^2$. That is, TransRate favors the diversity between different classes, while penalizes if the overlap between classes is high. Detailed proof and more theoretical analysis about the completeness and compactness can be found in Appendix D.3.

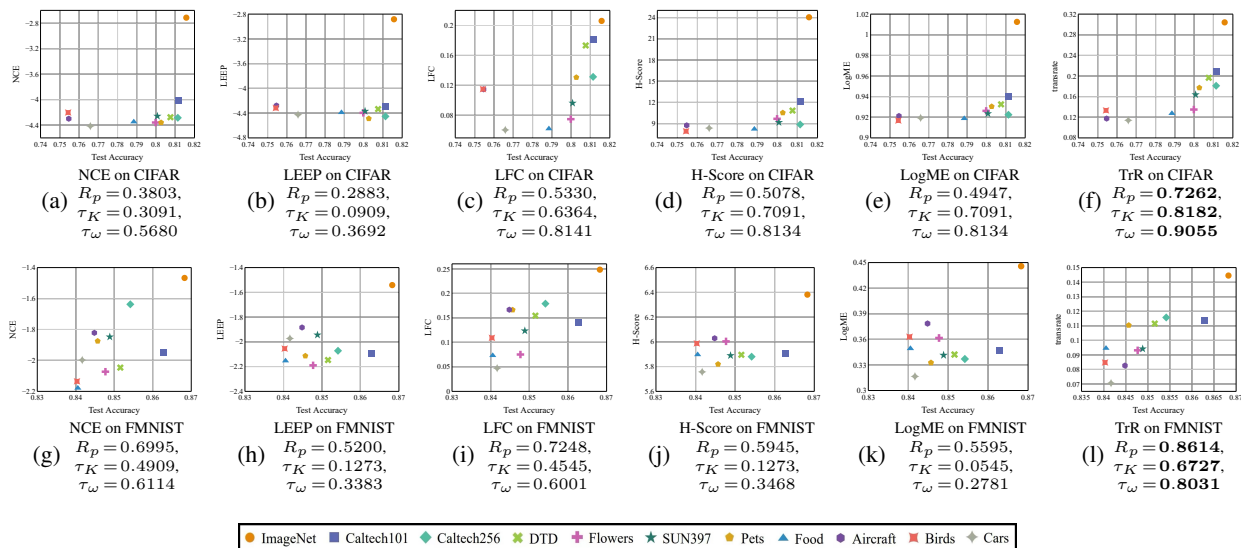


Figure 4: Transferability estimation on transferring ResNet-18 pre-trained on 11 different source datasets to CIFAR-100 and FMNIST.

4. Experiments

In this section, we evaluate the correlation between predicted transferability by TransRate and the transfer learning performance in various settings and for different tasks. Due to page limit, experiments covering more settings and the wall-clock time comparison are available in Appendix B.

4.1. Implementation Details

We consider fine-tuning a pre-trained model from a source dataset to the target task without access to any source data. For fine-tuning of the target task, the feature extractor is initialized by the pre-trained model. Then the feature extractor together with a randomly initialized head is optimized by running SGD on a cross-entropy loss for 100 epoches. The batch size (16, 32, 64, 128), learning rate (from 0.0001 to 0.1) and weight decay (from 1E-6 to 1E-4) are determined via grid search of the best average transfer performance over 10 runs on a validation set. The reported transfer performance is an average of top 5 accuracies over 20 runs of experiments under the best hyperparameters above.

Before performing fine-tuning on the target task, we calculate TransRate and the other baseline transferability measures on training examples of a target task. To compute the proposed TransRate score, we first run a single forward pass of the pre-trained model through all target examples to extract their features \hat{Z} , and then centralize \hat{Z} to have zero mean. Second, we compute the TransRate score as $R(\hat{Z}, \epsilon) - R(\hat{Z}, \epsilon|Y)$. In the experiments, we set $\epsilon = 1E-4$ by default. Since the scales of features extracted by different feature extractors may vary a lot, we scale the features by $1/\sqrt{\text{tr}(\hat{Z}\hat{Z}^\top)}$, such that the trace of the variance matrix of the normalized \hat{Z} is consistently equal to 1 for all models.

In the experiments of source selection and model selection, the features extracted by the pre-trained model trained on different source datasets or with different network architectures have significantly different patterns, making it difficult to directly compare their TransRate. To tackle this problem and improve the performance of TransRate, we project the variance matrix $\hat{Z}\hat{Z}^\top$ and $\hat{Z}^c\hat{Z}^{c\top}$ by a low-rank matrix $(\hat{Z}\hat{Z}^\top)^{-1}\hat{U}\hat{U}^\top$, where \hat{U} is a matrix whose c -th row is the centroid feature of c -th class.

We adopt LEEP (Nguyen et al., 2020), NCE (Tran et al., 2019), Label-Feature Correlation (LFC) (Deshpande et al., 2021), H-Score (Bao et al., 2019) and LogME (You et al., 2021) as the baseline methods. For a fair comparison, we assume no data from source tasks is available. In this scenario, the NCE score, defined by $-H(Y|Y_S)$ where Y_S is the labels from the source task, cannot be computed following the procedure described in its original paper. Instead, we follow (Nguyen et al., 2020) to replace Y_S with the softmax label generated by the classifier of a pre-trained model. Another setting for a fair comparison is that only one single forward pass through target examples is allowed for computational efficiency. In this case, we calculate H-score by pre-trained features and skip the computation of H-score based on the optimal target features as suggested in (Bao et al., 2019).

To measure the performance of TransRate and five baseline methods in estimating the transfer learning performance, we follow (Nguyen et al., 2020; Tran et al., 2019) to compute the Pearson correlation coefficient between the score and the average accuracy of the fine-tuned model on testing samples of the target set. The Kendall's τ (Kendall, 1938) and its variant, weighted τ , are also adopted as performance metrics. For brief, we denote the Pearson correlation coefficient, Kendall's τ and weighted τ by R_p , τ_K , and τ_ω , respectively.

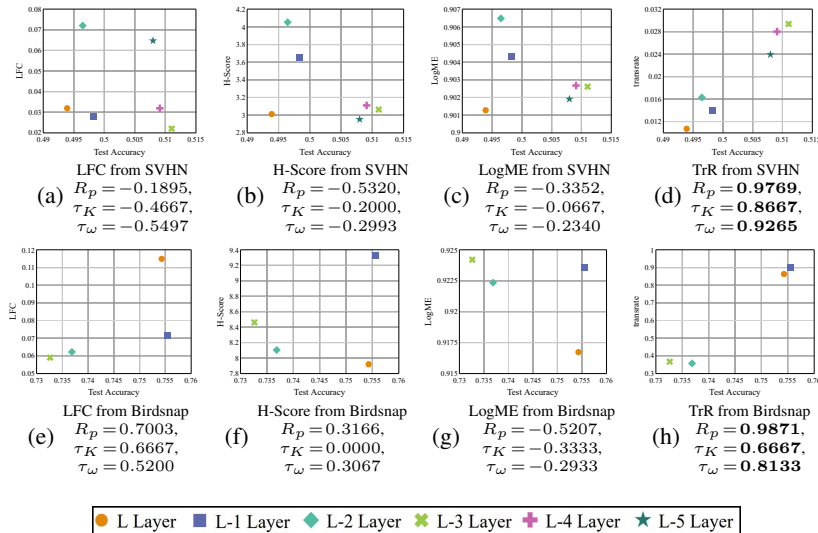


Figure 5: Transferability estimation on transferring different layers of ResNet-20 pre-trained on SVHN and of ResNet-18 pre-trained on Birdsnap to CIFAR-100.

4.2. Results

TransRate as a Criterion for Source Selection. One of the most important applications of TransRate is source model selection for a target task. Here we evaluate the performance of TransRate and other baseline measures for selection of a pre-trained model from 11 source datasets to a specific target task. The source datasets are ImageNet (Rusakovsky et al., 2015) and 10 image datasets from (Salman et al., 2020), including Caltech-101, Caltech-256, DTD, Flowers, SUN397, Pets, Food, Aircraft, Birds and Cars. For each source dataset, we pre-train a ResNet-18 (He et al., 2016), freeze it and discard the source data during fine-tuning. CIFAR-100 (Krizhevsky et al., 2009) and FMNIST (Xiao et al., 2017) are adopted as the target tasks. For all target datasets, we use the whole training set for fine-tuning and for transferability estimation. The details of these datasets and their pre-trained models are available in Appendix A; experiments on more target tasks are available in Appendix B.1.

Figure 4 show that LFC, H-Score, LogME and TransRate all correctly predict the ranking of top-5 source models, except the one pre-trained on Caltech-256, for CIFAR-100. TransRate achieves the best R_p , τ_K and τ_ω , which means that the ranking predicted by TransRate is more accurate than the others. As for FMNIST, TransRate correctly predicts the top-4 source models, though slightly underestimates the transferability of the Caltech-101 model, while all the other baselines fail to accurately predict the rank of the Caltech-101 model which comes the second among all. TransRate outperforms the baselines by a large margin in all correlation coefficients. These results demonstrate that TransRate can serve as a practical criterion for source selection in transfer

learning.

TransRate as a Criterion for Layer Selection. As introduced in Section 1, transferring different layers of the pre-trained model results in different accuracies; that is, the optimal layers to transfer are task-specific. To study the correlation between all transferability measures and the performance of transferring different layers, we conduct experiments on transferring the first layer to the K -th layer only. In this experiment, we consider transferring a ResNet-20 model pre-trained on SVHN or a ResNet-18 model pre-trained on Birdsnap to CIFAR-100. The candidate values of K for ResNet-20 and ResNet-18 are $\{9, 11, 13, 15, 17, 19\}$ and $\{11, 13, 15, 17\}$. The selected layers to transfer are initialized by the pre-trained model and the remaining ones are trained from scratch. NCE and LEEP are excluded in this experiment as they are not applicable to layer selection. For TransRate and other baselines, the transferability is estimated using the features extracted by the first K -th layer. Note that when K is not the last layer, we will apply the average pooling function on the features, which is used by the original ResNet in the last layer. More details about the experimental settings are available in Appendix A.

From Figure 5 we observe that TransRate is the only method that correctly predicts the layer with the highest performance in both experiments. In the experiment transferring different layers of the pre-trained model from SVHN, TransRate achieves the highest correlation coefficients. The baselines even have negative coefficients, which means that their predictions are inverse to the correct ranking. In the experiment transferring from Birdsnap, TransRate correctly predicts the rank of the top 2 layers with the highest transfer performance and also achieves the highest correlation coefficients.

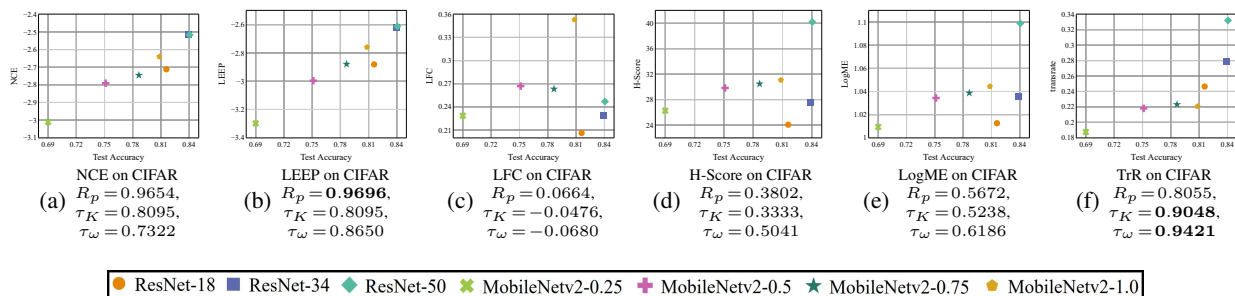


Figure 6: Result on transferring models with different architectures from ImageNet to CIFAR-100.

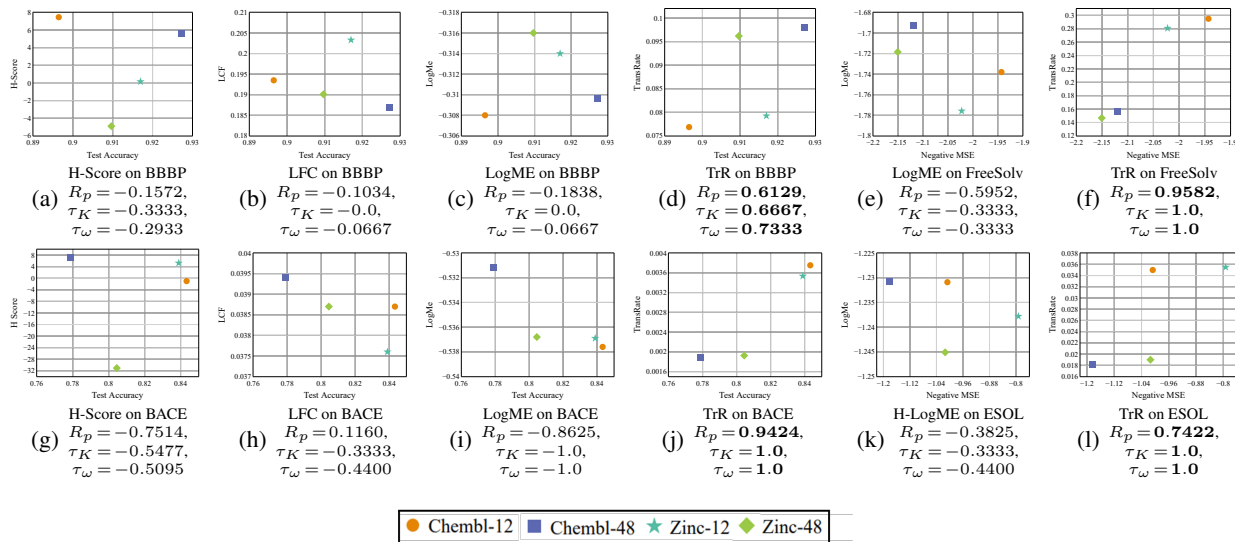


Figure 7: Result on transferring GNNs pre-trained on molecules sampled from different datasets to molecule property prediction tasks.

Both experiments demonstrate the superiority of TransRate in selecting the best layer for transfer. More experiments of layer selection with different source datasets, models, and target datasets are available in Appendix B.2.

TransRate as a Criterion for Pre-trained Model Selection.

Another important application of transferability measures is the selection of pre-trained models in different architectures. In practice, various pre-trained models on public large datasets are available. For example, PyTorch provides more than 20 pre-trained neural networks for ImageNet. To maximize the transfer performance from such a dataset, it is necessary to estimate the transferability of various model candidates and select the one with the maximal score to transfer. In this experiment, we consider seven kinds of pre-trained models on ImageNet to CIFAR-100. The seven types include ResNet18 (He et al., 2016), ResNet34, ResNet50, MobileNet0.25, MobileNet0.5, MobileNet0.75, MobileNet1.0 (Sandler et al., 2018). Figure 6 tells that TransRate in general has a significant linear correlation with the transfer accuracy, though it slightly underestimates MobileNet1.0. Though the predictions of LEEP and NCE achieve the best R_p , they rank ResNet-18 incorrectly with

underestimation. This also explains why they obtain lower τ_K and τ_ω than TransRate. The performances of LFC, H-Score and LogME are not as competitive as NCE, LEEP and TransRate, though. More experiments of model selection with more networks and target datasets are available in Appendix B.3.

Estimation of Unsupervised Pre-trained Models to Classification and Regression Tasks.

We also evaluate the effectiveness of TransRate and the baselines on estimating transferability from different unsupervised pre-trained models. The first type of self-supervised models we consider is GROVER (Rong et al., 2020) for graph neural networks (GNN). We evaluate the transferability of four candidate models by varying two types of architectures and two types of pre-trained datasets, which are denoted by ChemBL-12, ChemBL-48, Zinc-12, Zinc-48. We consider four target tasks that predict the molecular ADMET properties, including BBBP (Martins et al., 2012), BACE (Subramanian et al., 2016), Esol (Delaney, 2004), FreeSolv (Mobley & Guthrie, 2014). The BBBP and BACE are classification tasks, while Esol and FreeSolv are regression tasks. More details about the settings of the pre-trained GNN models and the datasets

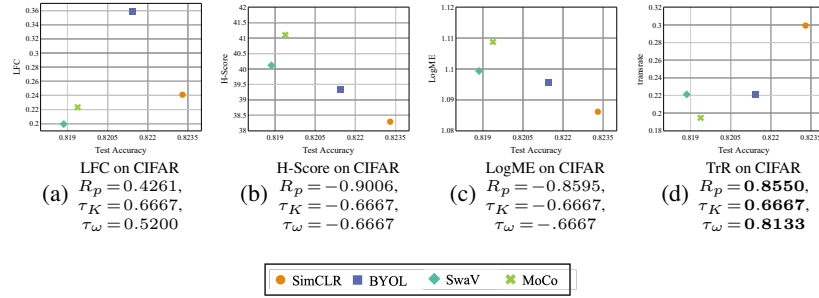


Figure 8: Result on transferring ResNet50 pre-trained with different self-supervised algorithms from ImageNet to CIFAR-100.

are available in Appendix A. Figure 7 show that in all four experiments, TransRate achieves the best performance regarding all 3 coefficients. To be specific, TransRate correctly predicts the ranking of all models in all experiments except the Zinc-48 model on BBBP. while the baselines all fail to predict the best model.

We also evaluate the performance in selecting 4 models pre-trained on ImageNet by 4 self-supervised algorithms, including SimCLR (Chen et al., 2020), BYOL (Grill et al., 2020), SwaV (Caron et al., 2020), MoCo (He et al., 2020). Figure 8 show that TransRate is the only method that correctly predicts the best-performing model. Though it overestimates the performance of SwaV, it still achieves the best correlation coefficients R_p , τ_K and τ_ω , outperforming the baseline methods by a large margin. Results on more target datasets are available in Appendix B.4. These results demonstrate the wide applicability as well as the effectiveness of TransRate in predicting the best unsupervised pre-trained model for regression or classification target tasks.

4.3. Discussion on Sensitivity to ϵ and Sample Size

As discussed in Figure 2, the approximation error in TransRate depends on 1) ϵ and 2) the sample size. By default, we have $\epsilon = 1E-4$, while we would investigate the sensitivity of TransRate to the value of ϵ . Appendix B.7 and Appendix D.4 demonstrate that as long as ϵ is less than a threshold, the performance of TransRate for estimating transferability and even the values of TransRate barely change.

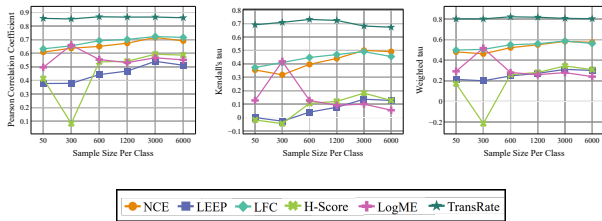


Figure 9: Influence of the sample size of target datasets to the performance of a transferability measure, when fine-tuning the pre-trained ResNet-18 from 11 different source datasets to FMNIST.

It is inevitable for both TransRate and the baselines to have the estimation error caused by a very limited number of samples that are insufficient to represent the true distribution. We would study the sensitivity of TransRate and the baseline methods with regard to the number of training samples in a target task. We adopt the same experiment settings as in source selection, except that the number of samples available for each class varies from 50 to 6000. The trends of the three types of correlation coefficients in Figure 9 speak that the performances of all algorithms generally drop when the sample size per class decreases. Unlike the baselines, the Kendall's τ and weighted τ of TransRate drop only by a minor percentage. This shows its superiority in predicting the correct ranking of the models, even when only a small number of samples are available for estimation.

5. Conclusion

In this paper, we propose a frustratingly easy transferability measure named TransRate that flexibly supports estimation for transferring both holistic and partial layers of a pre-trained model. TransRate estimates the mutual information between features extracted by a pre-trained model and labels with the coding rate. Both theoretic and empirical studies demonstrate that TransRate strongly correlates with the transfer learning performance, making it a qualified transferability measure for source dataset, model, and layer selection.

Acknowledgements

We would like to thank Yaodong Yu for the helpful discussion regarding the properties of coding rate. Ying Wei acknowledge the support of Project 9229073 by RMGS of Research Grants Council (RGC), Hong Kong.

References

Achille, A., Lam, M., Tewari, R., Ravichandran, A., Maji, S., Fowlkes, C. C., Soatto, S., and Perona, P. Task2vec: Task embedding for meta-learning. In *ICCV*, pp. 6430–6439, 2019.

- Agakov, D. B. F. The im algorithm: a variational approach to information maximization. *NeurIPS*, 16:201, 2004.
- Bao, Y., Li, Y., Huang, S.-L., Zhang, L., Zheng, L., Zamir, A., and Guibas, L. An information-theoretic approach to transferability in task transfer learning. In *ICIP*, pp. 2309–2313, 2019.
- Beirlant, J., Dudewicz, E. J., Györfi, L., and Van der Meulen, E. C. Nonparametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6(1):17–39, 1997.
- Belghazi, M. I., Baratin, A., Rajeshwar, S., Ozair, S., Bengio, Y., Courville, A., and Hjelm, D. Mutual information neural estimation. In *ICML*, pp. 531–540, 2018.
- Berg, T., Liu, J., Woo Lee, S., Alexander, M. L., Jacobs, D. W., and Belhumeur, P. N. Birdsnap: Large-scale fine-grained visual categorization of birds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2011–2018, 2014.
- Binia, J., Zakai, M., and Ziv, J. On the epsilon-entropy and the rate-distortion function of certain non-gaussian processes. *IEEE Transactions on Information Theory*, 20(4):517–524, 1974.
- Bossard, L., Guillaumin, M., and Van Gool, L. Food-101—mining discriminative components with random forests. In *European conference on computer vision*, pp. 446–461. Springer, 2014.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., and Joulin, A. Unsupervised learning of visual features by contrasting cluster assignments. *arXiv preprint arXiv:2006.09882*, 2020.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3606–3613, 2014.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Cui, Y., Song, Y., Sun, C., Howard, A., and Belongie, S. Large scale fine-grained categorization and domain-specific transfer learning. In *CVPR*, pp. 4109–4118, 2018.
- Delaney, J. S. Esol: estimating aqueous solubility directly from molecular structure. *Journal of chemical information and computer sciences*, 44(3):1000–1005, 2004.
- Deshpande, A., Achille, A., Ravichandran, A., Li, H., Zancato, L., Fowlkes, C., Bhotika, R., Soatto, S., and Perona, P. A linearized framework and a new benchmark for model selection for fine-tuning. *arXiv preprint arXiv:2102.00084*, 2021.
- Dwivedi, K. and Roig, G. Representation similarity analysis for efficient task taxonomy & transfer learning. In *CVPR*, pp. 12387–12396, 2019.
- Fei-Fei, L., Fergus, R., and Perona, P. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pp. 178–178. IEEE, 2004.
- Gaulton, A., Bellis, L. J., Bento, A. P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., et al. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic acids research*, 40(D1):D1100–D1107, 2012.
- Griffin, G., Holub, A., and Perona, P. Caltech-256 object category dataset. 2007.
- Grill, J.-B., Strub, F., Altché, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. A., Guo, Z. D., Azar, M. G., et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019.
- Kendall, M. G. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- Kraskov, A., Stögbauer, H., and Grassberger, P. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.
- Krause, J., Deng, J., Stark, M., and Fei-Fei, L. Collecting a large-scale dataset of fine-grained cars. 2013.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.

- Li, Y., Jia, X., Sang, R., Zhu, Y., Green, B., Wang, L., and Gong, B. Ranking neural checkpoints. *CVPR*, pp. 2663–2673, 2021.
- Ma, Y., Derksen, H., Hong, W., and Wright, J. Segmentation of multivariate mixed data via lossy data coding and compression. *IEEE transactions on pattern analysis and machine intelligence*, 29(9):1546–1562, 2007.
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., and Vedaldi, A. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Martins, I. F., Teixeira, A. L., Pinheiro, L., and Falcao, A. O. A bayesian approach to in silico blood-brain barrier penetration modeling. *Journal of chemical information and modeling*, 52(6):1686–1697, 2012.
- Mobley, D. L. and Guthrie, J. P. Freesolv: a database of experimental and calculated hydration free energies, with input files. *Journal of computer-aided molecular design*, 28(7):711–720, 2014.
- Moon, Y.-I., Rajagopalan, B., and Lall, U. Estimation of mutual information using kernel density estimators. *Physical Review E*, 52(3):2318, 1995.
- Nguyen, C. V., Hassner, T., Archambeau, C., and Seeger, M. Leep: A new measure to evaluate transferability of learned representations. *ICML*, 2020.
- Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729. IEEE, 2008.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Parkhi, O. M., Vedaldi, A., Zisserman, A., and Jawahar, C. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3498–3505. IEEE, 2012.
- Qin, Z., Kim, D., and Gedeon, T. Rethinking softmax with cross-entropy: Neural network classifier as mutual information estimator. *arXiv preprint arXiv:1911.10688*, 2019.
- Quinlan, J. R. Induction of decision trees. *Machine learning*, 1(1):81–106, 1986.
- Rong, Y., Bian, Y., Xu, T., Xie, W., Wei, Y., Huang, W., and Huang, J. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33, 2020.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Salman, H., Ilyas, A., Engstrom, L., Kapoor, A., and Madry, A. Do adversarially robust imagenet models transfer better? *arXiv preprint arXiv:2007.08489*, 2020.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, pp. 4510–4520, 2018.
- Shalev, Y., Painsky, A., and Ben-Gal, I. Neural joint entropy estimation. *arXiv preprint arXiv:2012.11197*, 2020.
- Song, J., Chen, Y., Ye, J., Wang, X., Shen, C., Mao, F., and Song, M. Depara: Deep attribution graph for deep knowledge transferability. In *CVPR*, pp. 3922–3930, 2020.
- Sterling, T. and Irwin, J. J. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.
- Subramanian, G., Ramsundar, B., Pande, V., and Denny, R. A. Computational modeling of β -secretase 1 (bace-1) inhibitors using ligand based approaches. *Journal of chemical information and modeling*, 56(10):1936–1949, 2016.
- Tishby, N. and Zaslavsky, N. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pp. 1–5. IEEE, 2015.
- Tong, X., Xu, X., Huang, S.-L., and Zheng, L. A mathematical framework for quantifying transferability in multi-source transfer learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Tran, A. T., Nguyen, C. V., and Hassner, T. Transferability and hardness of supervised classification tasks. In *ICCV*, pp. 1395–1405, 2019.
- Wang, Z., Dai, Z., Póczos, B., and Carbonell, J. Characterizing and avoiding negative transfer. In *CVPR*, pp. 11293–11302, 2019.
- Wu, Z., Ramsundar, B., Feinberg, E. N., Gomes, J., Geniesse, C., Pappu, A. S., Leswing, K., and Pande, V. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., and Torralba, A. Sun database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 3485–3492. IEEE, 2010.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *NeurIPS*, pp. 3320–3328, 2014.
- You, K., Liu, Y., Long, M., and Wang, J. Logme: Practical assessment of pre-trained models for transfer learning. *arXiv preprint arXiv:2102.11005*, 2021.
- Yu, Y., Chan, K. H. R., You, C., Song, C., and Ma, Y. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. *NeurIPS*, 33, 2020.
- Zamir, A. R., Sax, A., Shen, W., Guibas, L. J., Malik, J., and Savarese, S. Taskonomy: Disentangling task transfer learning. In *CVPR*, pp. 3712–3722, 2018.
- Zhang, G., Zhao, H., Yu, Y., and Poupart, P. Quantifying and improving transferability in domain generalization. *NeurIPS*, 2021.
- Zhang, W., Deng, L., and Wu, D. Overcoming negative transfer: A survey. *arXiv preprint arXiv:2009.00909*, 2020.

A. Omitted Experiment Details in Section 4

A.1. Image Datasets Description

Aircraft (Maji et al., 2013) The dataset consists of 10,000 aircraft images in 100 classes, with 66/34 or 67/33 training/testing images per class.

Birdsnap (Berg et al., 2014) The dataset has 49,829 images of 500 species of North American Birds. It is divided into a training set with 32,677 images and a testing set with 8,171 images.

Caltech-101 (Fei-Fei et al., 2004) The dataset contains 9,146 images from 101 object categories. The number of images in each category is between 40 and 800. Following (Salman et al., 2020), we sample 30 images per category as a training set and use the rest of images as a testing test.

Caltech-256 (Griffin et al., 2007) The dataset is an extension of Caltech 101. It contains 30,607 images spanning 257 categories. Each category has 80 to 800 images. We sample 60 images per category for training and use the rest of images for testing.

Cars (Krause et al., 2013) The dataset consists of 16,185 images of 196 classes of cars. It is divided into a training set with 8,144 images and a testing set with 8,041 images.

CIFAR-10 (Krizhevsky et al., 2009) The dataset contains 60,000 color images in 10 classes, with each image in the size of 32×32 . Each class has 5,000 training samples and 1,000 testing samples.

CIFAR-100 (Krizhevsky et al., 2009) The dataset is the same as CIFAR-10 except that it has 100 classes each of which contains 500 training images and 100 testing images.

DTD (Cimpoi et al., 2014) The dataset consists of 5,640 textural images with sizes ranging between 300×300 and 600×600 . There are a total of 47 categories, with 80 training and 40 testing images in each category.

Fashion MNIST (Xiao et al., 2017) The dataset involves 70,000 grayscale images from 10 classes with the size of each image as 28×28 . Each class has 6,000 training samples and 1,000 testing samples. Note that we limit the number of examples per class to be 300 when using Fashion MNIST as a target dataset.

Flowers (Nilsback & Zisserman, 2008) The dataset consists of 102 categories of flowers that are common in the United Kingdom. Each category contains between 40 and 258 images. We sample 20 images per category to construct the training set and use the rest of 6,149 images as the testing set.

Food (Bossard et al., 2014) The dataset contains 101,000 images organized by 101 types of food. Each type of food contains 750 training images and 250 testing images.

Pets (Parkhi et al., 2012) The dataset contains 7,049 images of pets belonging to 47 species. The training set contains 3,680 images and the testing set has 3,669 images.

SUN397 (Xiao et al., 2010) This dataset has 397 classes, each having 100 scenery pictures. For each class, there are 50 training samples and 50 testing samples.

A.2. Molecule Datasets Description

BBBP (Martins et al., 2012) The dataset contains 2,039 compounds with binary labels of blood-brain barriers penetration. In the dataset, 1560 compounds are positive and 479 compounds are negative.

BACE (Subramanian et al., 2016) The dataset involves 1,513 recording compounds which could act as the inhibitors of human β -secretase 1 (BACE-1). 691 samples are positive and the rest 824 samples are negative.

Esol (Delaney, 2004) The dataset documents the water solubility (log solubility in mols per litre) for 1,128 small molecules.

Freesolv (Mobley & Guthrie, 2014) The dataset contains 642 records of the hydration free energy of small molecules in water from both experiments and alchemical free energy calculations.

Note that on both BBBP and BACE we perform binary classification by training with a binary cross-entropy loss. Yet the regression tasks on Esol and Freesolv datasets are trained by an MSE loss. For all four datasets, we random split 80%

samples for training and 20% samples for testing.

A.3. Pre-trained Models

In the experiments, all models except GROVER used in Sec. 4.5 follow the standard architectures in PyTorch’s torchvision. For those models pre-trained on ImageNet, we directly download them from PyTorch’s torchvision. For those models pre-trained on other source datasets, we pre-train them with the hyperparameters obtained via grid search to guarantee the best performance. For GROVER (Rong et al., 2020), we run the released codes provided by the authors on two large-scale unlabelled molecule datasets. In total we have 4 pre-trained GROVER models, by considering two model architectures and pre-training on two datasets. The first model contains about 12 million parameters while the second has about 48 million parameters. Each of the two models is pre-trained on two unlabeled molecules datasets. The first one has 2 million molecules collected from ChemBL (Gaulton et al., 2012) and MoleculeNet (Wu et al., 2018); the second one contains 11 million molecules sampled from ZINC15 (Sterling & Irwin, 2015) and ChemBL. In total, we have 4 different pre-trained models, denoted by ChemBL-12, ChemBL-48, Zinc-12, Zinc-48.

A.4. Performance measure

We measure the performance of TransRaet and the baseline methods by Pearson correlation coefficient R_p , Kendall’s τ and weighted τ . The Pearson correlation coefficient measures the linear correlation between the predicted scores and transfer accuracies of the transfer tasks. Kendall’s τ , also known as Kendall rank correlation coefficient, measures the rank correlation, i.e. the similarity between two rankings ordered by transfer accuracy and transferability scores. Weighted τ is a variant of Kendall’s τ which focuses more on the top rankings. Since transferability scoring usually aims at selecting the best pre-trained model, we would highlight that weighted τ is the best one among these three measures for the transferability scoring method.

A.5. Details about the Layer Selection Experiments in Section 4.3

In the layer selection experiments, we transfer the first K layers from the pre-trained model and train the remaining layers from scratch. Notice that an average pooling function is applied to aggregate the outputs from different channels in the last layer of ResNet, which reduces the dimension of features. Inspired by this, we also apply an average pooling on the output of the K -th layer to reduce the feature dimension when we compute TransRate and other baseline transferability measures, even if K is not the last layer in layer selection. Besides the experiments on ResNet-20 and ResNet-18, we also conduct layer selection on ResNet-34 and present the results in Appendix B.2. We consider only the last layer of a block in ResNet as a candidate layer. The details of candidate layers and the feature dimensions are summarized in Table 2.

Table 2: The configurations of layer selection for three model architectures.

ResNet-20	Candidate Layers:	19	17	15	13	11	9		
	Feature Dimension:	64	64	64	32	32	32		
ResNet-18	Candidate Layers:	17	15	13	11				
	Feature Dimension:	512	512	256	256				
ResNet-34	Candidate Layers:	33	30	27	25	23	21	19	17
	Feature Dimension:	512	512	512	256	256	256	256	256

A.6. Details about Applying TransRate on Regression Tasks

When estimating TransRate on regression target tasks, the label (or target value) y_i is not discrete so that we cannot directly compute $R(\hat{Z}, \epsilon|Y)$ in TransRate. The key insight behind TransRate for classification is the overlap between the features of samples from different classes. Similarly, the transferability in a regression task can be estimated by the extent of the overlap between the features of samples with different target values. To realize this, we rank all n target values $\{y_i\}_{i=1}^n$ and divide them evenly into $C = 10$ ranges. For samples in each range, we compute $R(\hat{Z}^c, \epsilon)$ where \hat{Z}^c is the feature matrix of samples in the c -th range. Finally, we calculate $R(\hat{Z}, \epsilon|Y) = \sum_{c=1}^C R(\hat{Z}^c, \epsilon)$ and subtract $R(\hat{Z}, \epsilon|Y)$ from $R(\hat{Z}, \epsilon)$ to obtain the resulting TransRate.

A.7. Source Codes of TransRate

We implement the TransRate by Python. The codes are as follows:

```
import numpy as np

def coding_rate(Z, eps=1E-4):
    n, d = Z.shape
    (_, rate) = np.linalg.slogdet((np.eye(d) + 1 / (n * eps) * Z.transpose() @ Z))
    return 0.5 * rate

def transrate(Z, y, eps=1E-4):
    Z = Z - np.mean(Z, axis=0, keepdims=True)
    RZ = coding_rate(Z, eps)
    RZY = 0.
    K = int(y.max() + 1)
    for i in range(K):
        RZY += coding_rate(Z[(y == i).flatten()], eps)
    return RZ - RZY / K
```

Here, we observe that TransRate can be implemented by 10 lines of codes, which demonstrates its simplicity. It takes the features \hat{Z} (“Z” in the codes) and the labels Y (“y” in the codes), and distortion rate ϵ (“eps” in the codes) as input, calls the function “coding_rate” to calculate the coding rate of $R(\hat{Z}, \epsilon)$ and $R(\hat{Z}, \epsilon|Y)$, and finally returns the TransRate.

B. Extra Experiments

B.1. Extra Experiments on Source Selection

In Table 3, we summarize the results of source selection that we have presented in Section 4.2 and meanwhile include the results of source selection for 9 more target datasets. Since the 11 source datasets include the 9 new target datasets, we remove the model that is pre-trained on the target dataset and consider only the 10 models that are pre-trained on the remaining source datasets.

From Table 3, we observe that TransRate achieves 20/36 best performance and 7/36 second best performances, which indicates the superiority of TransRate across the spectrum of different target datasets. The most competitive baseline H-Score achieves 16/36 best performance. TransRate performs better in predicting the top transferability models, achieving 10/12 highest weighted τ and 7/12 highest Kendall’s τ , while H-Score correlates better linearly with the transfer accuracies, achieving 8/12 highest Pearson correlation coefficient. LogME has a similar performance to H-Score, but a little bit worse. As for NCE, LEEP, and LFC, their performances are not that satisfactory. These results indicate that TransRate serves as an excellent transferability predictor for selection of a pre-trained model from various source datasets.

B.2. Extra Results of Layer Selection

In this subsection, we provide more experiments on the selection of a layer for 15 models pre-trained on different source datasets, including the two experiments presented in Section 4.2. The settings are the same as those in Section 4.2, except that the pre-trained models are trained on different source datasets or with different architectures. Table 4 presents the results of the 15 experiments.

As shown in Table 4, TransRate correctly predicts the performance ranking of transferring different layers in 9 of 15 experiments, where its τ_K and τ_ω even equal to 1. In contrast, the best competitor, LFC, achieves all-correct prediction in only 3 experiments. This shows the superiority of TransRate over the three baseline methods in serving as a criterion for layer selection. In rare cases, H-Score and LogME achieve competitive or even better performances than TransRate, while they fail in most of the experiments. This hit-and-miss behavior can be explained by the assumption behind H-Score and LogME as mentioned in Section 2 – they consider the penultimate layer to be transferred only so that prediction of the transferability at other layers than the penultimate layer is not accurate.

Frustratingly Easy Transferability Estimation

Table 3: Transferability estimation on transferring ResNet-18 pre-trained on 11 different source datasets to different target datasets. The best performance among all transferability measures is highlighted in bold.

Target Datasets	Measures	NCE	LEEP	LFC	H-Score	LogME	TransRate
CIFAR-100	R_p	0.3803	0.2883	0.5330	0.5078	0.4947	0.7262
	τ_K	0.3091	0.0909	0.6364	0.7091	0.7091	0.8182
	τ_ω	0.5680	0.3692	0.8141	0.8134	0.8134	0.9055
FMNIST	R_p	0.6995	0.5200	0.7248	0.5945	0.5595	0.8614
	τ_K	0.4909	0.1273	0.4545	0.1273	0.0545	0.6727
	τ_ω	0.6114	0.3383	0.6001	0.3468	0.2781	0.8031
Aircraft	R_p	-0.8247	-0.5721	-0.6217	0.6758	0.6657	0.4722
	τ_K	-0.5111	-0.3333	-0.6000	0.5556	0.5556	0.5111
	τ_ω	-0.6180	-0.4054	-0.6854	0.5493	0.5493	0.5950
Birdsnap	R_p	0.3595	0.5460	0.7361	0.8682	0.8502	0.8677
	τ_K	0.2000	0.4667	0.2889	0.6444	0.6444	0.7333
	τ_ω	0.3995	0.5579	0.1115	0.5530	0.5530	0.6354
Caltech-101	R_p	0.2671	0.5255	0.5827	0.9058	0.8524	0.8962
	τ_K	0.1111	0.2444	0.4667	0.8667	0.8667	0.8667
	τ_ω	0.3578	0.4457	0.4801	0.8381	0.8381	0.8381
Caltech-256	R_p	0.3734	0.5655	0.5549	0.9080	0.8812	0.8600
	τ_K	0.2000	0.3333	0.3333	0.8667	0.8667	0.8222
	τ_ω	0.4953	0.5956	0.3415	0.8424	0.8424	0.9174
Cars	R_p	-0.6298	-0.1296	-0.0897	0.8385	0.8302	0.7356
	τ_K	-0.2444	-0.2889	-0.0667	0.7333	0.7333	0.6444
	τ_ω	-0.1193	0.0408	0.0375	0.8019	0.7273	0.7597
DTD	R_p	0.0218	0.1662	0.5243	0.9208	0.9293	0.9131
	τ_K	0.1556	0.2444	0.4222	0.6000	0.7333	0.7778
	τ_ω	0.1366	0.3519	0.3699	0.4409	0.7079	0.7755
Flowers	R_p	-0.3360	-0.2790	0.2631	0.8385	0.7967	0.9509
	τ_K	-0.2889	-0.2444	-0.0222	0.6000	0.6444	0.7778
	τ_ω	-0.1258	0.0865	-0.1650	0.5176	0.6035	0.7868
Food	R_p	0.2485	0.4300	0.5656	0.9243	0.9169	0.9065
	τ_K	0.1556	0.3333	0.2444	0.6000	0.6000	0.6000
	τ_ω	0.3214	0.4878	0.0860	0.4927	0.4927	0.5335
Pets	R_p	0.3512	0.4672	0.8306	0.9368	0.9019	0.8805
	τ_K	0.2444	0.5111	0.8667	0.8222	0.8222	0.7778
	τ_ω	0.3987	0.5679	0.8927	0.7351	0.7351	0.7148
SUN397	R_p	0.1535	0.4424	0.3693	0.9169	0.9058	0.7219
	τ_K	0.0222	0.2889	0.2000	0.7333	0.7333	0.5111
	τ_ω	0.3315	0.5159	0.1194	0.5928	0.5928	0.6424

B.3. Extra Results on Model Selection

In this subsection, we summarize experiment results on model selection for 8 target datasets, including the experiment for CIFAR-100 presented in Figure 6 and new experiments for 7 new target datasets.

As shown in Table 5, TransRate achieves 16/24 best performance and correctly predicts the transferability ranking of all models in 3 experiments (i.e., Caltech-101, Caltech-256 and Pets). NCE and LEEP both achieve 3/8 best Pearson correlation coefficient and LogME achieves 3/8 best weighted τ . However, in most experiments, their performances are not as competitive as TransRate.

We also conduct model selection experiments with 5 more candidate architectures, including DenseNet121, DenseNet169, DenseNet201, InceptionV3 and NASNet and present the results in Table 6. As more models are considered in the ranking, the model selection becomes more difficult. Compared to the result in Table 5, the performance in most experiments drops. Even so, TransRate still achieves 18/24 best performance, significantly outperforming the baseline measures.

Frustratingly Easy Transferability Estimation

Table 4: Transferability estimation on transferring different layers of the pre-trained model to the CIFAR-100 dataset. The best performance among all transferability measures is highlighted in bold.

	Measures	LFC	H-Score	LogME	TransRate
Source: SVHN Model: ResNet-20	R_p	-0.1895	-0.5320	-0.3352	0.9769
	τ_K	-0.4667	-0.2000	-0.0667	0.8667
	τ_ω	-0.5497	-0.2993	-0.2340	0.9265
Source: CIFAR-10 Model: ResNet-20	R_p	0.5755	0.6476	0.6551	0.6347
	τ_K	0.4667	0.2000	0.2000	0.4667
	τ_ω	0.4041	0.3673	0.3673	0.5224
Source: ImageNet Model: ResNet-18	R_p	0.2595	0.9876	0.9898	0.9866
	τ_K	0.0	1.0	1.0	1.0
	τ_ω	0.0	1.0	1.0	1.0
Source: ImageNet Model: ResNet-34	R_p	0.6997	0.9357	0.9370	0.9550
	τ_K	0.3333	0.9444	0.9444	0.9444
	τ_ω	0.4834	0.8674	0.8674	0.8674
Source: Aircraft Model: ResNet-18	R_p	0.6299	0.7983	0.0929	0.9560
	τ_K	0.3333	0.6667	0.0000	0.6667
	τ_ω	0.3333	0.8133	0.3067	0.8133
Source: Birdsnap Model: ResNet-18	R_p	0.7003	0.3166	-0.5207	0.9871
	τ_K	0.6667	0.6667	-0.3333	0.6667
	τ_ω	0.5200	0.3067	-0.2933	0.8133
Source: Caltech-101 Model: ResNet-18	R_p	0.9310	0.9015	0.8561	0.9871
	τ_K	1.0	0.6667	0.6667	1.0
	τ_ω	1.0	0.5200	0.5200	1.0
Source: Caltech-256 Model: ResNet-18	R_p	0.8395	0.1649	-0.4235	0.9763
	τ_K	0.6667	-0.3333	-0.3333	1.0
	τ_ω	0.8133	-0.2933	-0.2933	1.0
Source: Cars Model: ResNet-18	R_p	-0.2438	0.3188	-0.4489	0.9790
	τ_K	-0.3333	0.0000	-0.3333	0.6667
	τ_ω	-0.4400	0.3067	-0.2933	0.8133
Source: DTD Model: ResNet-18	R_p	0.9542	0.8818	0.7200	0.9860
	τ_K	1.0	0.6667	0.6667	1.0
	τ_ω	1.0	0.5200	0.5200	1.0
Source: Flowers Model: ResNet-18	R_p	0.8365	0.6054	0.0392	0.9925
	τ_K	0.3333	0.6667	0.0	1.0
	τ_ω	0.3333	0.5200	-0.0667	1.0
Source: Food Model: ResNet-18	R_p	0.7963	0.5642	-0.4487	0.8002
	τ_K	1.0	0.3333	-0.3333	1.0
	τ_ω	1.0	0.3333	-0.2933	1.0
Source: Pets Model: ResNet-18	R_p	0.7127	0.8349	0.6301	0.9608
	τ_K	0.6667	0.6667	0.3333	1.0
	τ_ω	0.8133	0.5200	0.2000	1.0
Source: SUN397 Model: ResNet-18	R_p	0.7647	0.6109	0.0409	0.9761
	τ_K	0.6667	0.6667	0.0	1.0
	τ_ω	0.8133	0.5200	-0.0667	1.0
Source: ImageNet Model: ResNet-18 Target: FMNIST	R_p	-0.0361	0.1775	0.1808	0.9169
	τ_K	-0.3333	0.0	0.0	1.0
	τ_ω	-0.4400	-0.1733	-0.1733	1.0

Frustratingly Easy Transferability Estimation

Table 5: Transferability estimation on transferring models with different architectures (ResNet18, ResNet34, ResNet50, MobileNet0.25, MobileNet0.5, MobileNet0.75 and MobileNet1.0) pre-trained on ImageNet to different target datasets. The best performance among all transferability measures is highlighted in bold.

Target Datasets	Measures	NCE	LEEP	LFC	H-Score	LogME	TransRate
CIFAR-100	R_p	0.9654	0.9696	0.0664	0.3802	0.5672	0.8055
	τ_K	0.8095	0.8095	-0.0476	0.3333	0.5238	0.9048
	τ_ω	0.7322	0.8650	-0.0680	0.5041	0.6186	0.9421
FMNIST	R_p	-0.5561	-0.4857	-0.4234	0.2182	0.1140	0.3649
	τ_K	-0.3333	-0.2381	-0.3333	0.3333	0.1429	0.4286
	τ_ω	-0.3088	-0.3751	-0.3581	0.2978	0.1515	0.4870
Aircraft	R_p	-0.4664	0.7383	0.6676	0.1350	0.6925	0.7952
	τ_K	-0.2000	0.4667	0.6000	0.3333	0.6000	0.7333
	τ_ω	-0.2639	0.4136	0.5184	0.5374	0.6857	0.6599
Caltech-101	R_p	0.9779	0.9748	0.5583	0.1241	0.7894	0.9648
	τ_K	0.8095	0.8095	0.3333	0.2381	0.8095	1.0
	τ_ω	0.7322	0.7322	0.2158	0.4345	0.8939	1.0
Caltech-256	R_p	0.9861	0.9851	0.5476	0.4262	0.6998	0.9626
	τ_K	0.8095	0.8095	0.4286	0.4286	0.7143	1.0
	τ_ω	0.7322	0.7322	0.3133	0.5937	0.7868	1.0
Cars	R_p	0.7317	0.9771	0.6308	0.1161	0.8319	0.7627
	τ_K	0.2381	0.8095	0.4286	0.3333	0.7143	0.8095
	τ_ω	0.2308	0.6786	0.3186	0.5416	0.8168	0.8125
Pets	R_p	0.9881	0.9867	0.9001	0.5085	0.8531	0.8643
	τ_K	0.8095	0.9048	0.7143	0.4286	1.0	1.0
	τ_ω	0.7322	0.9250	0.5286	0.5937	1.0	1.0
SUN397	R_p	0.9612	0.9638	0.5854	0.3117	0.7723	0.9609
	τ_K	0.7143	0.8095	0.5238	0.4286	0.7143	0.9048
	τ_ω	0.5983	0.6786	0.4633	0.6616	0.8168	0.8929

B.4. Extra Results on Self-supervised Model Selection

In this subsection, we conduct extra experiments on model selection among 4 self-supervised algorithms for 8 new target datasets and summarize their results and the experiment result presented in Figure 8 in Section 4.2.

The results are presented in Table 7. TransRate significantly outperforms the baseline methods. It achieves the best performance in all experiments except the experiment for Caltech-101. In the experiments for FMNIST, Caltech-256, Flowers and SUN397, TransRate correctly predict the ranking of all models, obtaining $\tau_K = 1$ and $\tau_\omega = 1$. The baseline methods all achieve only 3/27 best performance, but underperform in most experiments.

B.5. Extra Experiments on Sample Size Sensitivity Study

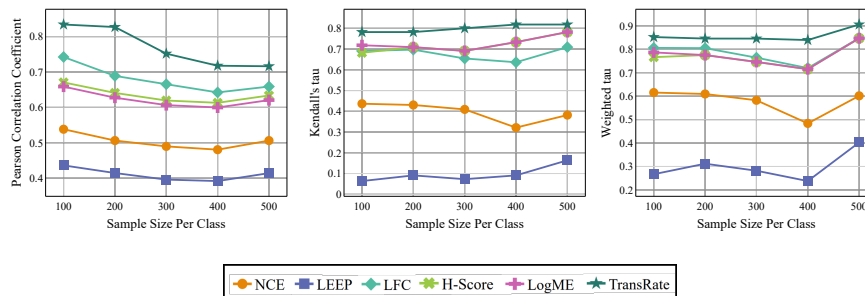


Figure 10: The three types of correlation coefficients between estimated transferability and test accuracy when varying the size of the target dataset. The correlation coefficient is measured by a series of transfer tasks that fine-tune the pre-trained ResNet-18 from 11 different source datasets to CIFAR-100.

We provide a supplementary experiment to further investigate the influence of sample size on the performance of transfer-

Frustratingly Easy Transferability Estimation

Table 6: Transferability estimation on transferring models with different architectures (ResNet18, ResNet34, ResNet50, MobileNet0.25, MobileNet0.5, MobileNet0.75, MobileNet1.0, *DenseNet121*, *DenseNet169*, *DenseNet201*, *InceptionV3* and *NASNet*) pre-trained on ImageNet to different target datasets. The best performance among all transferability measures is highlighted in bold.

Target Datasets	Measures	NCE	LEEP	LFC	H-Score	LogME	TransRate
CIFAR-100	R_p	0.7937	0.8506	-0.2159	0.5016	0.4965	0.8780
	τ_K	0.7436	0.7179	-0.0256	0.4872	0.4103	0.9231
	τ_ω	0.8315	0.8485	-0.0126	0.6058	0.5130	0.8498
FMNIST	R_p	0.2708	0.3522	0.3085	0.6226	0.1521	0.7086
	τ_K	0.0769	0.1795	0.1795	0.5897	0.0	0.7179
	τ_ω	0.2091	0.4351	0.2230	0.6670	-0.1171	0.8592
Aircraft	R_p	-0.4969	0.7260	-0.3718	0.3196	0.4179	0.6320
	τ_K	-0.3939	0.2727	0.2121	0.3939	0.4848	0.5758
	τ_ω	-0.3585	0.1134	0.0290	0.4971	0.6519	0.6838
Caltech-101	R_p	0.7955	0.8445	-0.1485	0.3112	0.5336	0.6392
	τ_K	0.6410	0.6154	-0.1026	0.5385	0.6923	0.7692
	τ_ω	0.6358	0.5380	-0.3313	0.7665	0.8214	0.8511
Caltech-256	R_p	0.9339	0.9199	0.3244	0.6125	0.6220	0.8100
	τ_K	0.7949	0.7179	0.1538	0.7179	0.7949	0.8974
	τ_ω	0.6922	0.6253	-0.0587	0.8668	0.8673	0.8971
Cars	R_p	0.4317	0.8114	-0.1935	0.2936	0.5289	0.7309
	τ_K	0.3590	0.7692	0.2821	0.4103	0.6154	0.7949
	τ_ω	0.5391	0.7974	0.2155	0.5620	0.6950	0.6894
Pets	R_p	0.9681	0.9787	0.6892	0.6333	0.7098	0.8143
	τ_K	0.7692	0.8205	0.5128	0.6154	0.8462	0.8718
	τ_ω	0.7178	0.7741	0.6214	0.6901	0.8439	0.8530
SUN397	R_p	0.9513	0.9166	0.3982	0.5834	0.7053	0.8380
	τ_K	0.7692	0.7692	0.3077	0.6667	0.7692	0.8462
	τ_ω	0.7332	0.7368	0.1686	0.7627	0.7577	0.7998

ability estimation algorithms. We vary the number of samples per class in CIFAR-100 from 100 to 500, and visualize the trend of the three types of correlation coefficients in Figure 10.

From Figure 10, we observe that the ranking prediction coefficients (i.e. τ_K and τ_ω) of all algorithms generally drops when the sample size per class decreases from 500 to 50. The deterioration of the performance is caused by the inaccurate estimation given a small number of samples. Even with slight deterioration given fewer samples, TransRate still outperforms the baselines. This shows its superiority in predicting the correct ranking of the models, even when only a small number of samples are available for estimation.

B.6. Time Complexity

In this subsection, we compare the running time of TransRate as well as the baselines. We run the experiments on a server with 12 Intel Xeon Platinum 8255C 2.50GHz CPU and a single P40 GPU. We consider three transfer tasks: 1) transferring ResNet-18 pre-trained on ImageNet to CIFAR-100 with full data; 2) transferring ResNet-18 pre-trained on ImageNet to CIFAR-100 with 1/10 data (50 samples per class); 3) transferring ResNet-50 pre-trained on ImageNet to CIFAR-100 with full data. For task 1), $n = 50,000$, $d = 512$; for task 2), $n = 5,000$, $d = 512$; for task 3), $n = 50,000$, $d = 2048$. The batch size in all experiments is 50.

We present the results in Table 8. First of all, we can observe that the time for fine-tuning a model is about 300 times of the time for transferability estimation (including the time for feature extraction and the time for computing a transferability measure). Besides, it often requires more than 10 times of fine-tuning to search for the best-performing hyper-parameters. Therefore, running a transferability estimation algorithm can achieve $3000\times$ speedup when selecting a pre-trained model and the layer of it to transfer. This highlights the necessity and importance of developing transferability estimation algorithms. Second, though LEEP and NCE computing the similarity between labels only show the highest efficiency, they suffer from the unsatisfactory performance in source selection and the incapability of accommodating unsupervised pre-trained models and layer selection. Third, amongst all the feature-based transferability measures, TransRate takes the shortest wall-clock time. This indicates its computational efficiency. The time costs of both H-Score and LogME are higher than TransRate,

Frustratingly Easy Transferability Estimation

Table 7: Transferability estimation on transferring models pre-trained with different self-supervised learning algorithms on ImageNet to different target datasets. The best performance among all transferability measures is highlighted in bold.

Target Datasets	Measures	LFC	H-Score	LogME	TransRate
CIFAR-100	R_p	0.4261	-0.9006	-0.8595	0.8550
	τ_K	0.6667	-0.6667	-0.6667	0.6667
	τ_ω	0.5200	-0.6667	-0.6667	0.8133
FMNIST	R_p	0.6259	0.2483	0.3307	0.9955
	τ_K	1.0	0.6667	0.6667	1.0
	τ_ω	1.0	0.8133	0.8133	1.0
Aircraft	R_p	-0.3821	0.2395	0.0404	0.9688
	τ_K	-0.3333	0.3333	0.0	0.6667
	τ_ω	-0.4400	0.2000	-0.1733	0.7333
Birdsnap	R_p	0.3656	0.5404	0.1623	0.6397
	τ_K	0.6667	0.3333	0.0	0.6667
	τ_ω	0.5200	0.3333	0.0	0.5200
Caltech-101	R_p	0.6620	0.7456	0.8963	-0.0979
	τ_K	0.3333	0.6667	0.6667	0.3333
	τ_ω	0.5333	0.8133	0.5200	0.5333
Caltech-256	R_p	0.2690	-0.5860	-0.4753	0.9239
	τ_K	0.3333	-0.3333	-0.3333	1.0
	τ_ω	0.2000	-0.2933	-0.2933	1.0
Cars	R_p	-0.4311	0.3379	0.1503	0.7498
	τ_K	-0.3333	0.3333	0.3333	0.3333
	τ_ω	-0.4400	0.2000	0.2000	0.5333
Flowers	R_p	0.3755	0.7177	0.3576	0.8077
	τ_K	-0.3333	0.0	0.0	1.0
	τ_ω	-0.4400	-0.1733	-0.0667	1.0
SUN397	R_p	-0.4770	0.4464	0.2504	0.8180
	τ_K	-0.3333	0.0	0.0	1.0
	τ_ω	-0.4400	0.0	0.0	1.0

which recognizes the necessity of developing an optimization-free estimation algorithm.

B.7. Sensitivity to Value of ϵ

To evaluate the influence of ϵ on TransRate, we conduct experiments on the toy case presented in Section 3.2 and on the layer selection with the same settings in Section 4.3. We vary the value of ϵ from 0.01 to 1E-15 and report the TransRate score and the performance (evaluated by R_p , τ , and τ_ω) of TransRate under different values of ϵ in Figure 11.

We have the following three observations. First, the TransRate scores hardly change when $\epsilon \leq 1E-3$ in the toy case and when $\epsilon \leq 1E-12$ in the layer selection experiment. This verifies the analysis Appendix D.4 that the value of TransRate does not change for a sufficiently small ϵ . Second, though the value of TransRate scores is still changing, their ranking does change for all ϵ in Figure 11(a) and for $\epsilon \leq 1E-3$ in Figure 11(b). Third, we see in Figure 11(c) that the performance of TransRate remains nearly the same when $\epsilon \leq 1E-3$. The second and third observations verify that the value of ϵ has limited influence to the performance of TransRate.

B.8. Target Selection

We follow (Nguyen et al., 2020) to also evaluate the correlation of the proposed TransRate and other baselines to the accuracy of transferring a pre-trained model to different target tasks. The target tasks are constructed by sampling different subsets of classes from a target dataset. We consider two target datasets: CIFAR-100 and FMNIST. For CIFAR-100, we construct the target tasks by sampling 2, 5, 10, 25, 50, and 100 classes. For FMNIST, we construct the target tasks by sampling 2, 4, 5, 6, 8, 10 classes. As shown in Proposition 1, the optimal log-likelihood is linear proportional to TransRate score minus the entropy of Y . In target selection, the entropy of Y can be different for different targets. Hence, we subtract $H(Y)$ from the TransRate in this experiment.

Table 8: Comparison of the computational cost of different measures.

	ResNet-18, Full Data		ResNet-18, Small Data		ResNet-50, Full Data	
	Wall-clock time (second)	Speedup	Wall-clock time (second)	Speedup	Wall-clock time (second)	Speedup
Fine-tune	8399.65	1×	882.33	1×	2.3×10^4	1×
Extract feature	30.1416		3.2986		72.787	
NCE	0.9126	9,204×	0.2119	4,164×	2.1220	10,839×
LEEP	0.7771	10,808×	0.1211	7,286×	1.9152	12,009×
LFC	30.1416	279×	0.7987	1,106×	149.3040	154×
H-Score	1.6285	5,158×	0.3998	2,207×	13.07	1,760×
LogME	9.2737	906×	2.0224	436×	50.1797	458×
TransRate	1.3410	6,264×	0.2697	3,272×	10.6498	2,160×

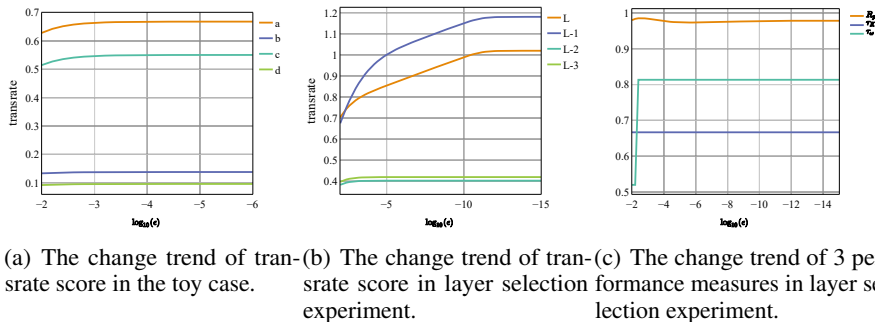


Figure 11: Sensitivity analysis of the value of ϵ . The figure in the left shows results on toy example introduced in section 3.2. The figures in middle and right columns show results on a layer selection experiment from a ResNet-18 model pre-trained on Birdsnap to CIFAR-100.

The results on transferring a pre-trained ResNet-18 on Imagenet to two target datasets and transferring a pre-trained ResNet-20 on CIFAR-10 to two target datasets are summarized in Table 9.

The results in Table 9 show that TransRate outshines the baselines in all 4 experiments. In both the first and third experiments, TransRate obtains the best performance in all three metrics. In the second experiment, it gets the best R_p , although its τ_K and τ_ω are a little bit lower than LFC. In the fourth experiment, TransRate achieves the best τ_K and τ_ω and the third best R_p (0.9449), which is competitive with the best R_p (0.9639). Generally speaking, TransRate is the best among all 6 transferability estimation algorithms. Yet we suggest that the competitors, NCE and LEEP, can be good substitutes if TransRate is not considered.

C. Theoretical studies of TransRate

C.1. Coding Rate and Shannon Entropy of a Quantized Continuous Random Variable

Rate-distortion function is known as ϵ -Entropy (Binia et al., 1974), which closely related to information entropy. As introduced in (Ma et al., 2007), coding rate is a regularizer version of rate-distortion function for the Gaussian source $\mathcal{N}(0, ZZ^\top)$. So coding rate is also closely related to information entropy. Even for non-Gaussian distribution, as presented in Appendix A in (Ma et al., 2007), the coding rate estimates a tight upper bound of the total number of nats needed to encode the region spanned by the vectors in Z subject to the mean squared error ϵ in each dimension. So it naturally related to the information in the quantized Z . We would note that the ϵ in (Ma et al., 2007) considers the overall distortion rate across all dimension while in our paper, the ϵ considers only one dimension. So the ϵ in our paper equals to ϵ^2/d in (Ma et al., 2007).

Notice that the code rate $R(Z, \epsilon)$ would be infinite if $\epsilon \rightarrow 0$. This property matches the property Shannon entropy of a quantized random variable that $H(Z^\Delta) = H(Z) - \log(\Delta)$ (Theorem 8.3.1 (Cover, 1999)) would be infinite when $\Delta \rightarrow 0$. Such a property does not hurt the estimation of mutual information as

$$MI(Z; Y) = H(Z) - H(Z|Y) = (H(Z) - \log(\Delta)) - (H(Z|Y) - \log(\Delta)) \approx H(Z^\Delta) - H(Z^\Delta|Y) \approx R(Z, \epsilon) - R(Z|Y, \epsilon),$$

Table 9: Transferability estimation on transferring from a pre-trained model to different target tasks.

Target Tasks	Measures	NCE	LEEP	LFC	H-Score	LogME	TransRate
Source: ImageNet	R_p	0.9810	0.9781	0.7599	-0.9597	-0.7940	0.9841
Model: ResNet-18	τ_K	0.9467	0.9467	0.6000	-0.9467	-0.6267	0.9733
Target: CIFAR-100	τ_ω	0.9537	0.9537	0.7662	-0.9600	-0.6104	0.9810
Source: ImageNet	R_p	0.9258	0.9246	0.8229	-0.8674	0.7200	0.9410
Model: ResNet-18	τ_K	0.7067	0.7067	0.8134	-0.7067	0.6533	0.7333
Target: FMNIST	τ_ω	0.7578	0.7578	0.8893	-0.7271	0.7782	0.8122
Source: CIFAR-10	R_p	0.9815	0.9819	0.5714	-0.9192	-0.9188	0.9896
Model: ResNet-20	τ_K	1.0	1.0	0.3600	-0.9733	-0.8133	1.0
Target: CIFAR-100	τ_ω	1.0	1.0	0.5320	-0.9695	-0.7878	1.0
Source: CIFAR-10	R_p	0.9622	0.9639	0.6410	-0.8871	0.6687	0.9449
Model: ResNet-20	τ_K	0.8400	0.8400	0.4667	-0.8133	0.3333	0.8400
Target: FMNIST	τ_ω	0.8579	0.8579	0.5262	-0.8267	0.5654	0.8797

where $\Delta = \sqrt{2\pi\epsilon\epsilon}$.

C.2. TransRate Score and Transfer Performance

In transfer learning, the pre-trained model is optimized by maximizing the log-likelihood, and the accuracy is closely related to the log-likelihood. As presented in Proposition 1, the ideal TransRate score highly aligns with the optimal log-likelihood of a task. This indicates that the ideal TransRate is closely related to the transfer performance.

Notice that the practical TransRate is not an exact estimation of the ideal TransRate. But generally, it is linearly proportional to the ideal TransRate. Together with Proposition 1 and the definition of transferability (in Definition 1), we get that the practical TransRate is larger than the transferability up to a multiplicative and/or an additive constant and smaller than the transferability up to another multiplicative and/or another additive constant. This means that:

$$\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon) \propto \text{Trf}_{T_s \rightarrow T_t}(g).$$

We also show in Lemma D.3 that the value of TransRate is related to the separability of the data from different classes. On one hand, TransRate achieves minimal value when the data covariance matrices of all classes are the same. In this case, it is impossible to separate the data from different classes and no classifier can perform better than random guesses. On the other hand, TransRate achieves its maximal value when the data from different classes are independent. In this case, there exists an optimal classifier that can correctly predict the labels of the data from different classes. The upper and lower bound of TransRate show that TransRate is related to the separability of the data, and thus, related to the performance of the optimal classifier.

D. Theoretical Details Omitted in Section 3

D.1. Proof of Proposition 1.

Proposition 1. Assume the target task has a uniform label distribution, i.e. $p(y = y^c) = \frac{1}{C}$ holds for all $c = 1, 2, \dots, C$. We then have:

$$\text{TrR}_{T_s \rightarrow T_t}(g) - H(Y) \gtrsim \mathcal{L}(g, h^*) \gtrsim \text{TrR}_{T_s \rightarrow T_t}(g) - H(Y) - H(Z^\Delta).$$

Proof. Firstly, we note that

$$\text{TrR}_{T_s \rightarrow T_t}(g) = H(Z) - H(Z|Y) = \text{MI}(Y; Z) \approx H(Z^\Delta) - H(Z^\Delta|Y) \tag{5}$$

As presented in (Agakov, 2004), the mutual information has a variational lower bound as $\text{MI}(Y; Z) \geq \mathbb{E}_{Z, Y} \log \frac{Q(z, y)}{P(z)P(y)}$ for a variational distribution Q . Following (Qin et al., 2019), we choose Q as

$$Q(z, y) = P(z)P(y) \frac{y \exp(h^*(z))}{\mathbb{E}_{y'} y' \exp(h^*(z))}, \tag{6}$$

where $h^*(z)$ is the output of optimal classifier before softmax, y is the one-hot label of z , y' is any possible one-hot label. If $p(y' = y^c) = \frac{1}{C}$ holds for all $c = 1, 2, \dots, C$, we have

$$\begin{aligned}
 \text{MI}(Y; Z) &\geq \mathbb{E}_{Z, Y} \log \frac{y \exp(h^*(z))}{\mathbb{E}_{y'} y' \exp(h^*(z))} \\
 &= \mathbb{E}_{Z, Y} \log \frac{y \exp(h^*(z))}{\frac{1}{C} \sum_{c=1}^C y^c \exp(h^*(z))} \\
 &= \mathbb{E}_{Z, Y} \log \frac{y \exp(h^*(z))}{\sum_{c=1}^C y^c \exp(h^*(z))} - \log\left(\frac{1}{C}\right) \\
 &\approx \mathcal{L}(g, h^*) + \log(C) \\
 &= \mathcal{L}(g, h^*) + H(Y)
 \end{aligned} \tag{7}$$

The last equality comes from the definition of the negative log-likelihood loss $\mathcal{L}(g, h^*)$, which is an empirical estimation of $\mathbb{E}_{Z, Y} \log \frac{y \exp(h^*(z))}{\sum_{c=1}^C y^c \exp(h^*(z))}$. Combining Eqs.(5) and (7), we have the first inequality in Proposition 1.

For proving the second inequality, we consider a classifier \bar{h} that predicts the label y for any data by $p(y) = \int_{z \in Z} p(y|z)p(z)dz$. The the empirical loss computed on this classifier is

$$\begin{aligned}
 \mathcal{L}(g, \bar{h}) &= \frac{1}{n} \sum_{i=1}^n \log p(y_i) = \frac{1}{n} \sum_{i=1}^n \log \left(\int_{z \in Z} p(y_i|z)p(z) \right) dz \\
 &\geq \frac{1}{n} \sum_{i=1}^n \log p(y_i|z_i) + \frac{1}{n} \sum_{i=1}^n \log p(z_i)
 \end{aligned} \tag{8}$$

The inequality comes from replacing the integral by one of its elements. It is easy to verify that the first term is an empirical estimation of $-H(Y|Z^\Delta)$ and the second term is an empirical estimation of $-H(Z^\Delta)$. By the definition of $\mathcal{L}(g, h^*)$, we have

$$\mathcal{L}(g, h^*) \geq \mathcal{L}(g, \bar{h}) \gtrsim -H(Y|Z^\Delta) - H(Z^\Delta) = \text{TrR}_{T_s \rightarrow T_t}(g) - H(Y) - H(Z^\Delta). \tag{9}$$

Combining Eqs. (7) and (9), we complete the proof. \square

Remark: Since the maximal log-likelihood is a variational form of mutual information between inputs and labels, the gap between the $MI(Y; Z)$ and the optimal log-likelihood is small. That is Eq. (7) is a tight upper bound of the optimal log-likelihood. The lower bound is proved by constructing a classifier \bar{h} without considering the feature. Such a classifier generally does not performs well in practice. The performance gap between the optimal classifier h^* and \bar{h} is larger. So the lower bound is loose. The lower bounds provided in NCE (Tran et al., 2019) and LEEP (Nguyen et al., 2020) are also proved through a similar technique. This means the lower bounds in our paper and in NCE, LEEP are all loose. But only TransRate is proved to be a tight upper bound of the maximal log-likelihood.

D.2. Properties of Coding Rate and TransRate Score

In this part, we discuss the properties of coding rate and TransRate Score.

Lemma D.1. For any $\hat{Z} \in \mathbb{R}^{d \times n}$, we have

$$R(\hat{Z}, \epsilon) = \frac{1}{2} \log \det \left(I_d + \frac{1}{n\epsilon} \hat{Z} \hat{Z}^\top \right) = \frac{1}{2} \log \det \left(I_n + \frac{1}{n\epsilon} \hat{Z}^\top \hat{Z} \right)$$

Lemma D.1 presents the commutative property of the coding rate, which is known in (Ma et al., 2007). Based on this lemma, we can reduce the complexity of $\log \det(\cdot)$ computation in $R(\hat{Z}, \epsilon)$ from $\mathcal{O}(d^{2.373})$ to $\mathcal{O}(n^{2.373})$ if $n < d$.

Lemma D.2. For any $\hat{Z} \in \mathbb{R}^{d \times n}$ having r singular values, denoted by $\sigma_1, \sigma_2, \dots, \sigma_r$, we have

$$R(\hat{Z}, \epsilon) = \frac{1}{2} \sum_{i=1}^r \log \left(1 + \frac{1}{n\epsilon} \sigma_i^2 \right)$$

Lemma D.2 is an inference from the invariant property of the coding rate presented in (Ma et al., 2007). Both lemmas are inferred from Sylvester's determinant theorem.

Lemma D.3. (*Upper and lower bounds of TransRate*) For any $\hat{Z}^c \in \mathbb{R}^{d \times n_c}$ for $c = 1, 2, \dots, C$, let $\hat{Z} = [\hat{Z}^1, \hat{Z}^2, \dots, \hat{Z}^C]$ which is the concatenation of all \hat{Z}^c . We then have

$$\text{TrR}_{T_s \rightarrow T_t}(g, \epsilon) = R(\hat{Z}, \epsilon) - \sum_{c=1}^C \frac{n_c}{n} R(\hat{Z}^c, \epsilon) \geq 0$$

The equality holds when $\frac{\hat{Z}\hat{Z}^\top}{n} = \frac{\hat{Z}^c(\hat{Z}^c)^\top}{n_c}$ for all c .

$$\begin{aligned} \text{TrR}_{T_s \rightarrow T_t}(g, \epsilon) &= R(\hat{Z}, \epsilon) - \sum_{c=1}^C \frac{n_c}{n} R(\hat{Z}^c, \epsilon) \\ &\leq \frac{1}{2} \sum_{c=1}^C \left(\log \det \left(I_d + \frac{1}{n\epsilon} \hat{Z}^c(\hat{Z}^c)^\top \right) - \frac{n_c}{n} \log \det \left(I_d + \frac{1}{n_c\epsilon} \hat{Z}^c(\hat{Z}^c)^\top \right) \right) \end{aligned}$$

The equality holds when $\hat{Z}^{c_1}(\hat{Z}^{c_2})^\top = 0$ for all $1 \leq c_1 < c_2 \leq C$.

The proof follows the upper and lower bound of $R(\hat{Z}, \epsilon)$ in Lemma A.4 of (Yu et al., 2020).

D.3. Proof of the toy case in Section 3.2

In the end of Section 3.2, we present a toy case of a binary classification problem with $\hat{Z} = [\hat{Z}^1, \hat{Z}^2]$ where $\hat{Z}^1 \in \mathbb{R}^{d \times n/2}$ and $\hat{Z}^2 \in \mathbb{R}^{d \times n/2}$. The lower and upper bound of the TransRate in this toy case can be derived by Lemma D.3. But in Sec. 3.2, we derived the bounds in another way. Here, we provide details of the derivation.

By Lemma D.1, we have

$$\begin{aligned} R(\hat{Z}, \epsilon) &= \frac{1}{2} \log \det(I_{2n} + \alpha \hat{Z}^\top \hat{Z}) \\ &= \frac{1}{2} \log \det \left(I_n + \alpha \begin{bmatrix} (\hat{Z}^1)^\top \hat{Z}^1 & (\hat{Z}^1)^\top \hat{Z}^2 \\ (\hat{Z}^2)^\top \hat{Z}^1 & (\hat{Z}^2)^\top \hat{Z}^2 \end{bmatrix} \right) \\ &= \frac{1}{2} \log \det \left\{ (I_{n/2} + \alpha (\hat{Z}^1)^\top \hat{Z}^1 + \alpha (\hat{Z}^2)^\top \hat{Z}^2) \right. \\ &\quad \left. + \alpha \left[(\hat{Z}^1)^\top \hat{Z}^1 (\hat{Z}^2)^\top \hat{Z}^2 - (\hat{Z}^1)^\top \hat{Z}^2 (\hat{Z}^2)^\top \hat{Z}^1 \right] \right\} \\ &\geq \frac{1}{2} \log \det \left\{ (I_{n/2} + \alpha (\hat{Z}^1)^\top \hat{Z}^1 + \alpha (\hat{Z}^2)^\top \hat{Z}^2) \right\} \\ &\quad + \frac{\alpha}{2} \log \det \left\{ (\hat{Z}^1)^\top \hat{Z}^1 (\hat{Z}^2)^\top \hat{Z}^2 - (\hat{Z}^1)^\top \hat{Z}^2 (\hat{Z}^2)^\top \hat{Z}^1 \right\} \end{aligned} \quad (10)$$

The first equality comes from Lemma D.1; the third equality follows the property of matrix determinant that for any square matrices A, B, C, D with the same size, i.e.,

$$\det \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \det(AD - BC).$$

The inequality follows that for any positive definite symmetric matrices A and B , $\det(A + B) \geq \det(A) + \det(B)$.

From Eqn. (10), we know that when $(\hat{Z}^1)^\top \hat{Z}^1$ and $(\hat{Z}^2)^\top \hat{Z}^2$ are fixed, the lower bound of TransRate is related to $\log \det \left\{ \left[(\hat{Z}^1)^\top \hat{Z}^1 (\hat{Z}^2)^\top \hat{Z}^2 - (\hat{Z}^1)^\top \hat{Z}^2 (\hat{Z}^2)^\top \hat{Z}^1 \right] \right\}$. The value of this term is determined by $(\hat{Z}^1)^\top \hat{Z}^2$. So the value of TransRate is also determined by $(\hat{Z}^1)^\top \hat{Z}^2$, i.e., the overlap between the two classes. When \hat{Z}^1 and \hat{Z}^2 are completely the same, this term becomes zero and TransRate achieves its minimal value as $\frac{1}{2} \log \det \left\{ (I_{n/2} + \alpha (\hat{Z}^1)^\top \hat{Z}^1 + \alpha (\hat{Z}^2)^\top \hat{Z}^2) \right\}$. When \hat{Z}^1 and \hat{Z}^2 are independent, this term achieves its maximal value as $\log \det \left((\hat{Z}^1)^\top \hat{Z}^1 (\hat{Z}^2)^\top \hat{Z}^2 \right)$ and TransRate achieves its maximal value as well.

D.4. The Influence of ϵ

In Section 3.2, we mention that the value of ϵ has minimal influence on the performance of TransRate. Here we provide more details about the influence of the choice of ϵ on the value of TransRate.

Assume that we scale ϵ by a positive scaler α . After scaling, the value of $R(\hat{Z}, \alpha\epsilon)$ and $R(\hat{Z}, \alpha\epsilon|Y)$ is different from $R(\hat{Z}, \epsilon)$ and $R(\hat{Z}, \epsilon|Y)$. By Lemma D.2, we have

$$R(\hat{Z}, \alpha\epsilon) = \frac{1}{2} \sum_{i=1}^r \log\left(1 + \frac{1}{n\alpha\epsilon} \sigma_i^2\right).$$

If $\frac{1}{n\alpha\epsilon} \sigma_i^2 \gg 1$, which holds for a sufficiently small ϵ , we have $\log\left(1 + \frac{1}{n\alpha\epsilon} \sigma_i^2\right) \approx \log\left(\frac{1}{n\alpha\epsilon} \sigma_i^2\right) = \log\left(\frac{1}{n\epsilon} \sigma_i^2\right) - 2\log(\alpha)$. Then for sufficiently small ϵ , α and sufficiently large σ_i , we have $R(\hat{Z}, \alpha\epsilon) \approx R(\hat{Z}, \epsilon) - r\log(\alpha)$ and $R(\hat{Z}, \alpha\epsilon|Y) \approx R(\hat{Z}, \epsilon|Y) - r\log(\alpha)$. Therefore, the influence of α is nearly canceled in calculating TransRate as we subtract $R(\hat{Z}, \alpha\epsilon)$ by $R(\hat{Z}, \alpha\epsilon|Y)$. Though this assumption may not hold in practice, we further verify the influence empirically in Appendix B.7.

D.5. Time Complexity

When $d < n$ and $d < n_c$, the computational costs of $R(\hat{Z}, \epsilon)$ and $R(\hat{Z}, \epsilon|Y)$ are $\mathcal{O}(d^{2.373} + nd^2)$ and $\mathcal{O}(Cd^{2.373} + Cn_c d^2)$, respectively. Thus, the total computation cost of TransRate is $\mathcal{O}((C+1)d^{2.373} + 2nd^2)$. By Lemma D.1, when $n < d$ or $n_c < d$, we can further reduce the cost of computing $R(\hat{Z}, \epsilon)$ or $R(\hat{Z}, \epsilon|Y)$. Besides, we can implement the computation of $R(\hat{Z}, \epsilon)$ and $R(\hat{Z}^c, \epsilon)$ in parallel, so that we can reduce the computation cost of TransRate to $\mathcal{O}(\min\{d^{2.373} + nd^2, n^{2.373} + dn^2\})$.