Discriminability-Transferability Trade-Off: An Information-Theoretic Perspective

Quan Cui^{1,2†}, Bingchen Zhao^{1,3†}, Zhao-Min Chen^{1,4}, Borui Zhao¹ Renjie Song¹⊠, Boyan Zhou⁵, Jiajun Liang¹, and Osamu Yoshie²

 1 MEGVII Technology 2 Waseda University 3 University of Edinburgh 4 Wenzhou University 5 ByteDance

cui-quan@toki.waseda.jp, zhaobc.gm@gmail.com, songrenjie@megvii.com

Abstract. This work simultaneously considers the discriminability and transferability properties of deep representations in the typical supervised learning task, *i.e.*, image classification. By a comprehensive temporal analysis, we observe a trade-off between these two properties. The discriminability keeps increasing with the training progressing while the transferability intensely diminishes in the later training period. From the perspective of information-bottleneck theory, we reveal that the incompatibility between discriminability and transferability is attributed to the over-compression of input information. More importantly, we investigate why and how the InfoNCE loss can alleviate the over-compression, and further present a learning framework, named contrastive temporal coding (CTC), to counteract the over-compression and alleviate the incompatibility. Extensive experiments validate that CTC successfully mitigates the incompatibility, yielding discriminative and transferable representations. Noticeable improvements are achieved on the image classification task and challenging transfer learning tasks. We hope that this work will raise the significance of the transferability property in the conventional supervised learning setting.

Keywords: Information-Bottleneck Theory, Representation Learning, Discriminability, Transferability, Contrastive Learning

1 Introduction

In recent decades, great progress has been achieved on learning discriminative representations with deep neural networks in a supervised learning manner. Advanced by such powerful deep representations, performances of many real-world computer vision applications are remarkably improved, *e.g.*, visual categorization[17, 20, 37], object detection [16, 34, 46], and semantic segmentation [28, 54]. However, the mainstream supervised learning works concentrate on pursuing more discriminative representations, which deserves the most attention for direct effects on model performances. Except for the discriminability, transferability

^{\dagger} denotes equal contributions, and \boxtimes denotes the corresponding author.



Fig. 1. Considered properties of deep representations. In a vanilla training, we reveal that the discriminability keeps increasing while the transferability first climbs to a peak and then gradually decreases. To address this issue, we propose the contrastive temporal coding, which successfully alleviates the incompatibility between transferability and discriminability.

should also be taken into considerations in the conventional supervised learning classification task, which is preferred for many downstream tasks [4, 14, 49, 58, 57] yet neglected in the conventional supervised learning setting.

We start this work by investigating the correlation between discriminability and transferability properties of representations in the entire supervised training process. Concretely, in the training process of a deep classification model, we extract representations from each training epoch and respectively assess their discriminability and transferability. Conclusions are generally illustrated in the left part of Figure 1. We observe that the discriminability keeps getting better, while the transferability intensely diminishes in the later training. It reveals that representations can hardly be discriminative and transferable at the same time in the conventional supervised learning setting, *i.e.*, these two properties could be incompatible. Nevertheless, high-quality representations are expected to possess both properties, and we suppose that the learning mechanism underlying the conventional supervised learning leads to the trade-off.

Interestingly, we notice a concept Information-Bottleneck (IB) trade-off in the IB theory [1, 39, 42], *i.e.*, the network learning can be interpreted by finding the optimal trade-off between input information compression and label-related information enhancement. This IB trade-off intrigues us to explain our observed trade-off between discriminability and transferability. Following IB, we visualize mutual information dynamics [2] in the information plane [39], and reveal the overcompression phenomenon, *i.e.*, prolonged input information compression leads to inadequate information on downstream tasks and thus poor transferability.

The above IB-based perspective also provides us an idea to make the discriminability and transferability compatible, *i.e.*, simultaneously training the model with a loss for the specific task and another loss to *counteract the overcompression*. To support this standpoint, we establish the correlation between counteracting over-compression and improving transferability via a principle components analysis (PCA) perspective. We further explore and provide a solution based on the InfoNCE [32, 40, 19] loss to counteract the over-compression. Concretely, we present a two-stage learning framework, namely Contrastive Temporal Coding (CTC). The learning process of CTC consists of two stages, *i.e.*, the information aggregation and revitalization stages. In the first stage, a classification model is optimized and the last epoch model is stored as the information bank, which aggregates informative representations. In the second stage, we introduce an InfoNCE loss between the current model and the information bank, counteracting the over-compression. As shown in the right part of Figure 1, our proposed method successfully alleviates the incompatibility, achieving high-quality representations in supervised learning.

2 Related Work

Information-Bottleneck Theory. The information-bottleneck (IB) theory [1, 39, 42] provides an information-theoretic principle for encoding the input data into a compressed representation. The theory is based on measuring the mutual information between the input/label variable and the representation variable. It is demonstrated that the representation learning of deep neural networks undergoes two phases, *i.e.*, the empirical error minimization and representation compression phases. In the first phase, the mutual information on the label variable is rapidly increased. When it comes to the compression phase, most of the optimization epochs are spent on decreasing the mutual information on the input variable. From the perspective of IB theory, we analyze the incompatibility between the transferability and the discriminability of learned representations in the common supervised learning, which is mostly ignored by the community. We conjecture and prove that the drop of transferability owes to the prolonged compression phase, *i.e.*, the information relevant to downstream tasks is overly discarded for learning discriminative representations.

Discriminability vs. Transferability. The transferability of deep representations has been studied from various perspectives [23, 36, 11, 43, 52, 24, 3, 7]. And many previous works also noticed the correlation between discrimiability and transferability in the domain adaptation area [3, 7]. In this work, we venture to study the correlation between discriminability and transferability with from an information-theoretic point of view. It is worth noting that [23] also noticed that better ImageNet classification results (obtained from better loss functions) could lead to worse transfer learning performances, strongly supporting our work.

Contrastive Learning. The main idea of contrastive learning is maximizing the similarity between samples from the same category/view while minimizing the similarity between samples from different categories/views. In recent years, contrastive learning has shown great potentials in self-supervised and unsupervised learning [4, 14, 22, 32, 33, 40, 41, 49, 4, 5, 50, 57] but is not widely studied in the supervised learning setting [22, 38]. Different from these previous works, our method contrasts the representations from the current training against those from the previous training. The contrastive learning objective has been proven to be the lower bound of the mutual information between the two views [32, 40, 48]. In this work, we utilize this property for optimizing the mutual information between deep representations from different training stages.

4 Q. Cui and B. Zhao et al.

3 Discriminability vs. Transferability

In this section, we first reveal the incompatibility between *discriminability* and *transferability* properties of deep representations. Then, we give theoretical and empirical explanations from the information-theoretic perspective.

3.1 Revealing the Incompatibility

We optimize a classification model with the cross-entropy loss on CIFAR-100 [25] and evaluate these properties²:

Discriminability. Intuitively, the discriminability of representations is reflected by the classification accuracy. However, the high accuracy only indicates separable representations since a sample can be correctly classified but near located to the decision boundary. Previous work [37, 47] pointed out that the discriminability can be better revealed by nearest neighbor search (NNS) algorithms. Besides, a subtle but essential component in a classification model should be considered, *i.e.*, the classifier. Recent works [56, 21] revealed that the performance of a classification model is closely related to the quality of its classifier. To precisely quantify the discriminability, we propose to evaluate retrieval and clustering performances of representations, which are built upon classifier-irrelevant NNS algorithms.

To measure discriminability, we evaluate representations with a typical network architecture ResNet18 [17]. Representations from each training epoch are extracted and assessed by retrieval and clustering tasks. Recall@1 (R@1) is the evaluation metric for retrieval tasks. Normalized Mutual Information (NMI) is reported to assess clustering performances. Metrics are calculated only by the test set to avoid problems that can be caused by over-fitting the training set.

Transferability. For measuring the transferability, it is reasonable to transfer learned knowledge to out-of-sample datasets. Given a deep classification model learned on a source dataset, we freeze its backbone network and re-train a classifier on top of the last feature layer on unseen target datasets. Corresponding classification accuracies reveal the transferability.

To measure the transferability, CIFAR-100 is the source dataset, and target datasets are CINIC10 [10] and STL-10 [8], respectively. Following the above experimental settings, ResNet-18 model is utilized.

Trade-off between discriminability and transferability. As illustrated in Figure 2, it can be observed that both R@1 and NMI have been improved with the training process. However, the continual training in later epochs can significantly hurt the transferability. In all the above experiments, models with the best transferability are mostly located in the middle training stage, rather than the later stage, where the best discriminability is scored. Conclusively, with the transferability remains uncertain, and, more importantly, these two properties could be incompatible in the later training stage.

² Representations refer to the outputs of the backbone, which are processed with a global average pooling in popular models [17].



Fig. 2. Temporal analyses of representations in the vanilla training process. In each subfigure, the X-axis is the training epoch, and Y-axis is the evaluated metric (*e.g.*, R@1, NMI, and Top-1 accuracy). " \uparrow " denotes "the higher the better". In the entire training, it can be observed that the discriminability keeps becoming better while the transferability first goes to a peak and then intensely drops in the later training. In a typical supervised training process, discriminability and transferability are incompatible.

3.2 Connection to Information-Bottleneck Trade-Off

The trade-off between discriminability and transferability in Sec. 3.1 motivates us to explain it with the well-known Information-Bottleneck (IB) trade-off [42, 39, 1].

Recap of the IB theory. The IB theory explains the learning of deep neural networks (DNNs) by the Information-Bottleneck trade-off, *i.e.*, the network learning is interpreted by finding the optimal trade-off between input information compression and label-related information enhancement.

An essential viewpoint of IB theory is that, except for the input variable Xand label variable Y, the hidden representation layer T is regarded as a variable. Under these assumptions, the mutual information (MI) between X (or Y) and T is used to describe the trade-off between input information compression and label-related information enhancement. I(X;T) denotes the MI between inputs X and representations T, and I(T;Y) represents the MI between representations T and labels Y. Based on MI, the empirical error minimization (ERM) and representation compression phases are defined. The fast ERM phase rapidly increases I(T;Y), and subsequently, the much longer representation compression phase results in the decrease of I(X;T). These two phases indicate that network first rapidly memorizes label-related information, and then compresses input information for finding an optimal trade-off.

IB trade-off meets discriminability-transferability trade-off. In the IB trade-off, we suppose that input information could be *overly compressed* on the source dataset. Due to the network focuses on enhancing label-related information, the task-irrelevant information could be discarded. Thus, insufficient information on target datasets brings about the unsatisfactory transferability. In other words, the over-compression results in the aforementioned empirical observation of discriminability-transferability trade-off. To prove this standpoint, we follow IB to calculate MI dynamics on both source and target datasets, capturing the

6 Q. Cui and B. Zhao et al.

correlation between the information-bottleneck trade-off and discriminability-transferability trade-off 3 .

(1) MIs on the source dataset. As shown in Figure 3, expectedly, I(T;Y)on the source dataset (CIFAR-100) rapidly increases in initial training, corresponding to the fast ERM phase. Then, I(X;T) slowly decreases with the training progressing, matching the representation compression phase. Associated with results in Figure 2, the discriminability greatly benefits from the representation compression phase.



(2) MIs on the target dataset. Both Fig. 3. Mutual information on both source I(X;T) and I(T;Y) on target datasets and target datasets.

are illustrated in Figure 3. With representation compression on the source dataset, I(T; Y) on both STL-10 and CINIC-10 climbs to a peak and then gradually decreases, consistently acting like the transferring performances in Figure 2. Meanwhile, I(X;T) on target datasets also gradually decrease with the training progressing, indicating that the information relevant to target datasets is also discarded due to the long representation compression phase.

Over-compression leads to the incompatibility. The above results reveal the *over-compression* issue, which degrades both I(X;T) and I(T;Y) on target datasets in the later training. Driven by the IB trade-off, the transferability can be sacrificed by the supervised leaning to overly compress input-related information for efficient representations. Over-compression results in insufficient information related to target datasets, unsatisfactory transferring performances, and the incompatibility. However, we suppose that high-quality representations are expected to possess both properties, arising our next explorations.

4 Method

4.1 Alleviating the Incompatibility

Understanding the transferability. Following [29], we hypothesize that the transferability between a source dataset and a target dataset relies on the model learns important representation patterns shared by source and target datasets. Thus, the decrease of transferability is attributed to the information of representations which related to the target dataset is overly compressed.

Counteracting over-compression improves transferability. To achieve this goal, an intuitive solution is to counteract the decrease of I(X;T) on the source dataset. To support our motivation, we give an explanation from the Principal Components Analysis (PCA) perspective.

³ We use the Mutual Information Neural Estimation (MINE) [2] method to calculate the mutual information between continuous variables.

Given an example where the model is linear, maximizing I(X;T) is equivalent to solving a PCA on the source dataset⁴. Solving PCA helps the model capture the most representative representations on the source dataset. If the target dataset shares important patterns with the source dataset, maximizing I(X;T)on the source dataset ensures the most important representation patterns are captured, further improving transferring performances. Since calculating the mutual information for continuous variables is challenging, we introduce an approximation in the next. The definition of I(X;T) is as followed:

$$I(X;T) = H(T) - H(T|X),$$

where H(T) is the entropy of the latent representation T and H(T|X) is the conditional entropy. Since the support set of X contains tremendous natural images from the P(X, Y) and the neural network is deterministic, the conditional entropy equals 0 [31]. Thus, the above equation can be re-written as:

$$I(X;T) = H(T),$$

which demonstrates that counteracting the decreasing of H(T) on the source dataset could be a potential way to alleviate over-compression. In the following, we demonstrate how to counteract the decrease of H(T). It is proven in [19, 32, 40] that minimizing the InfoNCE loss maximizes a lower bound on mutual information. Given two representation variables T_1 and T_2 , the relation between InfoNCE loss and the mutual information $I(T_1; T_2)$ can be derived as:

$$I(T_1; T_2) \ge \log(N) - \mathcal{L}_{\text{InfoNCE}},\tag{1}$$

where N is normally the number of samples in a training set. As demonstrated, minimizing the InfoNCE loss would increase the lower bound of $I(T_1; T_2)$. Combining Eqn.(1) with the definition of the mutual information:

$$I(T_1; T_2) = H(T_1) + H(T_2) - H(T_1, T_2)$$

$$\leq H(T_1) + H(T_2) - \max(H(T_1), H(T_2)),$$

$$= \min(H(T_1), H(T_2)),$$

the Eqn.(1) could be re-written as:

$$\log(N) - \mathcal{L}_{\text{InfoNCE}} \le \min(H(T_1), H(T_2)).$$
(2)

Eqn.(2) suggests that minimizing InfoNCE loss improves the lower bound of $\min(H(T_1), H(T_2))$, which simultaneously improves lower bounds of $H(T_1)$ and $H(T_2)$. Therefore, we could select the representation T_1 from early training and fix it as constant, regard the representation in later training as T_2 , and develop an InfoNCE loss between the constant representation T_1 and the later representation T_2 , for counteracting the decrease of $H(T_2)$. Inspired by the well-known "memory bank" concept, we name the constant representation T_1 as the information bank.

Concretely, in Eqn.(2), if $H(T_2) \leq H(T_1)$, the InfoNCE loss would improve the lower bound of $H(T_2)$. If $H(T_2) > H(T_1)$, the objective of counteracting has been reached. In this manner, the loss $\mathcal{L}_{\text{InfoNCE}}$ ensures the representation of later training has a relatively large H(T).

⁴ Proofs are attached in the appendix A.1.

8 Q. Cui and B. Zhao et al.

4.2 Contrastive Temporal Coding

Inspired by the above explorations, we propose a two-stage learning framework to alleviate the incompatibility between discriminability and transferability. In the first training stage, named information aggregation stage, the main objective is to obtain the information bank to provide the $H(T_1)$ in Eqn.(2). In the second training stage, named information revitalization stage, the main objective is to further counteract the decrease of $H(T_2)$ via Eqn.(2).

Information aggregation stage (IAS). In this stage, a classification model is trained with a vanilla cross-entropy (CE) loss \mathcal{L}_{CE} as the main loss function for learning T_1 . Eqn.(2) indicates that, if $H(T_2) \leq H(T_1)$, the InfoNCE loss will improve the lower bound of $H(T_2)$. Thus, it is reasonable to choose the T_1 with a relatively large $H(T_1)$ in this stage. However, experiments in Sec. 3 demonstrate that, only with the CE loss, the information compression is not controllable and predictable, and the last epoch model could still be overly compressed. For practical usage, it is tricky to select a good T_1 among all epochs, and we hope the last model is a satisfactory choice. To this end, we introduce an auxiliary InfoNCE loss \mathcal{L}_{IAS} for desensitizing the model selection.

Let T_1 denote the representation variable. Similar to IB theory, we regard a memory bank V as a variable, which keeps representations for each training sample and provide contrastive samples. For $t_1 \in T_1$ and $v \in V$ (which are normalized representations), contrastive representation pairs are composed as $\{t_1^i, v^j\}_{i,j=1}^N$, and N is the number of training samples. Positive pairs are constructed by representations of the identical samples, while negative pairs are composed of different samples. Therefore, the loss function of t_1^i can be re-written as:

$$\mathcal{L}_{\text{IAS}} = -\log \frac{\exp(t_1^i \cdot v^i)}{\sum_{j=1}^N \exp(t_1^i \cdot v^j)},\tag{3}$$

where v^i is the positive "key". When the learning process ends, the model of the latest epoch is saved as the information bank \hat{T}_1 . Information revitalization stage (IRS). Let T_2 denote representation variable of the second stage. In this stage, we continue training the model from the end of the former stage, and develop an InfoNCE loss \mathcal{L}_{IRS} to counteract the decrease of $H(T_2)$. For $\hat{t}_1 \in \hat{T}_1$ and $t_2 \in T_2$, contrastive representation pairs are composed as $\{\hat{t}_1^i, t_2^j\}_{i,j=1}^N$, where t_2 is from the current epoch and \hat{t}_1 is from the previously saved information bank. The optimization objective function of t_2^j can be written as:

$$\mathcal{L}_{\rm IRS} = -\log \frac{\exp(t_2^j \cdot \hat{t}_1^j)}{\sum_{k=1}^N \exp(t_2^j \cdot \hat{t}_1^k)},\tag{4}$$

where \hat{t}_1^{γ} is the positive "key". Consequently, the optimization objective can implicitly counteract the decrease of $H(T_2)$ and over-compression, promoting the transferability of learned representations.

Learning framework. The learning of the information aggregation stage is driven by a cross-entropy loss function \mathcal{L}_{CE} and a contrastive loss \mathcal{L}_{IAS} as:

$$\mathcal{L}_{\text{stage 1}} = \alpha \mathcal{L}_{\text{IAS}} + \mathcal{L}_{\text{CE}}.$$
 (5)

At the end of this stage, the information bank model is saved. In the information revitalization stage, the learning of the current model is simultaneously motivated by a contrastive learning loss \mathcal{L}_{IRS} and a cross-entropy loss \mathcal{L}_{CE} . The \mathcal{L}_{IRS} is calculated with representations from both the current model and the information bank. The loss can be written as:

$$\mathcal{L}_{\text{stage 2}} = \beta \mathcal{L}_{\text{IRS}} + \mathcal{L}_{\text{CE}}.$$
 (6)

 α and β are weighting hyper-parameters, which respectively emphasize the importance of two properties. The general framework could be similar to the well-known knowledge distillation process [18, 13, 53]; however, learning the first stage without the \mathcal{L}_{IAS} loss function (which corresponds to the teacher model training) still leads to poor transferability in T_1 . Consequently, developing the \mathcal{L}_{IPS} loss with such T_1 will limit the increasing of T_2 .

4.3 Discussions on Self-Supervised Learning (SSL)

We notice that SSL has shown great potential in transfer learning tasks [49, 14, 6], while linear probing results of these methods still cannot catch up with supervised learning results. It is similar to supervised learning that the transferability and discriminability also cannot be compatible. However, the superiority of SSL in improving transferability drives us to explore the underlying working mechanism. **On transferability.** Mainstream SSL methods depend on InfoNCE loss, and we attempt to formulate its learning target from the information view. In representative works, a dictionary (memory bank [49] or momentum encoder [14, 6]) is typically used for providing negative samples for a trainable encoder. We denote the dictionary representation as T_1 , and the trainable encoder representation as T_2 . Given that T_1 comes from the dictionary, the $H(T_1)$ could be regarded as a constant before being updated. Minimizing the contrastive loss also counteracts the decrease of $H(T_2)$. Thus, the good transferability of SSL methods could also result from the learning target for informative representations.

On discriminability. Mainstream SSL methods can be interpreted by a (K+1)-way softmax-based classification task, where K is the number of negative samples [14]. Intuitively, due to K is a large number, the (K+1)-way classification task is more challenging than the conventional supervised image classification. Thus, we conjecture that SSL compresses input information slowly and is inferior to enhance label-related information. We also notice in [6] that longer training will bring in significant discriminability improvement but minor transferability improvement, which suggest that SSL methods could require longer time for compressing input information and enhancing label-related information.

Later experiments will demonstrate that our CTC outperforms SSL methods on the transferability property.

5 Experiments

Experiments and discussions focus on the following two parts: (1) Validating our motivation and (2) Transferring representations learned by our method.



Fig. 4. Temporal analysis of representations in the vanilla and CTC training processes. Results of the vanilla and CTC are colored in blue and red, respectively. Two stages are divided by grey dotted lines. Compared with the vanilla training, it can be observed that the transferability is greatly preserved by CTC in later epochs. Furthermore, discriminability is improved by preserving the transferability.

5.1 Motivation Validation and Main Results

Motivation: counteracting over-compression. In this part, we prove our motivation that the proposed method is able to counteract the over-compression. Following the experimental settings in

Sec. 3, we report a temporal analysis of our proposed method. Generally, ResNet18 [17] is trained with CTC on the CIFAR-100 [25] dataset, and representations of each epoch are evaluated and reported afterward. For discriminability, representations of models from each epoch are extracted to perform retrieval and clustering tasks. As to the transferability, we transfer learned representations to STL-10 [8] and CINIC-10 [10].

Evaluations of discriminability and transferability are illustrated in Figure 4. We divide two stages with a grey dotted line. In the first stage, benefitting from \mathcal{L}_{IAS} , our method (red lines) achieves bet-



Fig. 5. Mutual information of CTC on both source and target datasets. For better illustration, two stages are split with grey dotted lines. The left and right parts correspond to information aggregation and revitalization stages, respectively.

ter transferring results in the last epoch than the baseline (blue lines). The second stage further improves transferring results on target datasets (STL-10 and CINIC-10) by large margins. Mutual information dynamics are provided in Figure 5. Similar to transferring results, the I(T; Y) on target datasets also keeps increasing in the second stage, showing the information relevant to target datasets is revitalized by our CTC. Meanwhile, the I(X; T) on source and target datasets is also non-decreasing in the second stage. The above results jointly demonstrate that the over-compression has been successfully alleviated by our method. Besides, the discriminability is not damaged, proving the incompatibility is also mitigated.

Benchmarking on CIFAR-100 and ImageNet. To show the direct benefits of counteracting the over-compression, we benchmark our proposed CTC on

CIFAR-100 [25] and ImageNet [35]. For CIFAR-100, results are reported in Table 1. Our proposed method consistently outperforms the baseline method. For ImageNet, results are reported in Table 2 and better results are also achieved compared with the vanilla training results. Counteracting the over-compression does not damage the discriminability and conversely benefits the classification performances. Moreover, our CTC brings in no increase on the number of model parameters and FLOPs.

method	top-1 acc. $(\%)$	method	top-1 acc. $(\%)$	
Res18+CosLr	$79.3 {\pm} 0.2$	Res50+CosLr	76.1 ± 0.1	
Res18+CTC(Ours)	80.1 ± 0.3	Res50+CTC(Ours)	76.4 ± 0.1	
Table 1. Top-1 accuracies (%) on CIFAR-		Table 2. Top-1 accuracies (%) on Ima		
100 of baseline and CTC (5 runs).		geNet of baseline and CTC (3 runs).		

Boosting transferability with CTC. To further validate that (1) the correctness of the discriminability-transferability trade-off and (2) our method is able to adjust the trade-off, we conduct experiment on sacrificing the discriminability for boosting the transferability. Since the information aggregation stage of CTC decides the lower bound of H(T), we could adjust the hyper-parameter α for helping the model learn informative representations. Unavoidably, large weight for \mathcal{L}_{IAS} influences the learning of \mathcal{L}_{CE} . As shown in Table 6 and Figure 7, increasing α to 0.5 leads to a normal classification accuracy on CIFAR-100, but significantly better transferring results. It also suggests that ending the baseline at an appropriate time is a bad option for enhancing transferability. Moreover, early ending baseline would unavoidably lead to poor discriminability.

method	top-1 acc. (%)
Res18+CosLr	79.3 ± 0.2
$\text{Res}18 + \text{CTC}(\alpha = 0.1)$	$80.1 {\pm} 0.3$
$\mathrm{Res18}{+}\mathrm{CTC}(\alpha=0.5)$	79.4 ± 0.3

Fig. 6. Top-1 accuracies (%) on CIFAR-100. We adjust the α from 0.1 to 0.5, and the top-1 accuracy is decreased but still better than vanilla training. Transferring results are in the right figure.



Fig. 7. Setting CTC's α to 0.5 (red) leads to greatly better transferring results than vanilla training (blue), showing the overcompression is sufficiently alleviated.

5.2 Towards Better Transferability

In this part, we study (1) how to get better transferability of representations and (2) how to plug in CTC to further boost the transferability, proving its scalability. **AutoAugment (AA) contributes to good transferability.** We suppose that strong augmentations could contribute to good transferability, since strong augmentation can increase the difficulty of information compression, and thus lead to higher H(T). Except for normal crop and resize operations, color normalization and translation operations might be more difficult to learn. Thus, on the CIFAR-100 dataset, we first conduct temporal analysis on representations learned with



Fig. 8. Temporal analyses of representations in the **vanilla** and **CTC** training with the AutoAugment. Results of the vanilla and CTC are colored in blue and red, respectively. Two stages of CTC (red lines) are separated by the grey dotted line. Compared with the vanilla training (blue lines), it can be observed that the transferability of CTC keeps climbing in the entire learning process. Different to Figure 4, we train the baseline and CTC with the same number of epochs for fair comparisons.

method	top-1 acc. $(\%)$
Res18+AA+CosLr	$80.4 {\pm} 0.3$
Res18+AA+CTC(Ours)	81.2±0.2

urs)	81.2 ± 0.2	Res50+AA+CTC(Ours)
acies	(%) on CIFAR-	Table 4. Top-1 accuract

method

Res50+AA+CosLr

Table 3. Top-1 accuracies (%) on CIFAR-100 of baseline and CTC with AA.

Table 4. Top-1 accuracies (%) on ImageNet of baseline and CTC with AA.

 $\frac{\text{top-1 acc. (\%)}}{76.8\pm0.1}$

77.2±0.1

the AA [9], and results are blue lines in Figure 8. Similar to results from Sec. 3, the discriminability of representations becomes better with the training progressing. However, it is surprised to notice that the transferability of representations learned without AA, especially in later training epochs. It conveys an intuition that strong augmentations could contribute to good transferability, which possibly result from our hypothesis that strong augmentation can increase the difficulty of information compression.

"AA+CTC" contributes to better transferability. Subsequently, we prove our CTC is scalable and orthogonal to other methods. We first optimize the model with CTC and AA, and then report the temporal analysis by red lines in Figure 8. With the assistance of CTC, discriminability scores higher within expectations. Moreover, the transferability keeps climbing in the entire learning process, eventually achieving the best transferring classification accuracies which can not be achieved only with AA. It further validates our motivation that the transferability could be promoted by counteracting the decrease of representation entropy. For comprehensiveness, we benchmark our CTC with AutoAugment on CIFAR-100 and ImageNet. All experimental settings in this part are the same as those in Sec. 5.1. As reported in Table 3 and 4, superior classification accuracies than the baseline also meet our expectations.

pro training method	Performance					
pre-training method	AP^{bbox}	AP_{50}^{bbox}	AP_{75}^{bbox}	AP^{mask}	AP_{50}^{mask}	AP_{75}^{mask}
Res50 random init.	30.2	48.9	32.7	28.6	46.6	30.7
Res50+MoCo v2 [6]	38.5	58.3	41.6	33.6	54.8	35.6
Res50+InfoMin [41]	39.0	58.5	42.0	34.1	55.2	36.3
Res50+CosLr	38.2	58.2	41.2	33.3	54.7	35.2
Res50+CTC(Ours)	39.5	58.7	42.0	34.2	55.4	36.2

Table 5. COCO object detection and instance segmentation based on Mask-RCNN-FPN with $1 \times$ schedule.

5.3 Transferring Representations

In this part, we present extensive tasks and datasets to which representations learned by our method can be transferred. Since self-supervised learning methods exceed supervised learning in transferring tasks, we also compare with representative self-supervised learning works [14, 6, 41].

Object detection and instance segmentation. In this part, we transfer the learned representations of CTC to the object detection and instance segmentation tasks [15, 14, 46]. Models pre-trained on ImageNet are further fine-tuned with the Mask-RCNN-FPN [16, 26] on the MS-COCO dataset [27] in the commonly applied training protocol.

Transferring results are summarized in Table 5. Compared to the vanilla pre-training with cosine learning rate (CosLr), our CTC yields consistently better transfer performances on the COCO dataset. Compared with self-supervised learning methods, *i.e.*, MoCo v2 [6] and InfoMin [41], CTC also achieves comparable or better results on COCO. For one thing, it indicates that our method also preserves information concerning the detection task. For another, the superiority of SSL in transferring tasks is challenged, *i.e.*, our method proves that supervised pre-training has the potential to achieve better results than SSL methods.

pro training method	CUB200	Aircraft
pre-training method	top-1 acc. (%)	top-1 acc. $(\%)$
Res50+CosLr†	62.5	27.8
Res50+CTC(Ours)†	63.7	28.2
Res50+AA+CosLr [†]	64.8	31.2
$\frac{\text{Res50+AA+CTC(Ours)}}{1000}$	66.1	32.1
Res50+CosLr‡	80.1	82.5
Res50+CTC(Ours)‡	81.7	84.1
Res50+AA+CosLr‡	81.3	83.4
${\rm Res50{+}AA{+}CTC(Ours)}\ddagger$	83.5	85.6

Table 6. Top-1 classification accuracies (%) on transferring representations to FGVC datasets. "†" and "‡" denote the backbone network is frozen and unfrozen, respectively.

Fine-grained visual categorization. In this part, representations are transferred to popular fine-grained visual categorization (FGVC) [55, 12] datasets. With the backbone network pre-trained on ImageNet frozen and unfrozen, two kinds of transferring experiments are conducted on the CUB-200-2011 (CUB200) [45] and the FGVC-Aircraft (Aircraft) [30] datasets. To compare with SSL methods, we extend the experiment on the large-scale iNaturalist-18 (iNat-18) [44] dataset. Training details are attached in appendix.

Transferring results of CUB200 and Aircraft are reported in Table 6 and 7. By freezing the backbone parameters and training a linear classifier on top of the learned representations, we observe that the representations learned by CTC achieve better performance over the vanilla baseline. Results of training the model with unfrozen backbone are also provided in Table 6. CTC again outperforms the vanilla training, further validating its effectiveness and practicality. Notably, on iNat-18, our method also outperforms the MoCo v1 [14], even the model is pre-trained on the billion-level data Instagram-1B [51]. It further demonstrates that learning transferable representations in the supervised learning is a promising research direction.

pre-training method	iNat-18		
pro training motifou	top-1 acc. $(\%)$		
Res50+CosLr	66.1		
Res50+MoCo v1 (IN-1M) [14]	65.6		
Res50+MoCo v1 (IG-1B) [14]	65.8		
Res50+CTC (ours)	66.4		
Res50+AA+CosLr	66.3		
Res50+AA+CTC (ours)	66.7		

Table 7. Top-1 accuracies (%) on iNaturalist and the backbone is frozen. "IN-1M" and "IG-1B" denote pre-training with ImageNet-1M [35] and web Instagram-1B [51] datasets, respectively. All methods (except MoCo v1 IG-1B) are pre-trained on IN-1M.

6 Discussion and Conclusion

This study focuses on learning representations with good discriminability and transferability at the same time. The trade-off between these properties is firstly observed by us via a temporal analysis. To explain this incompatibility, we explore the correlation between information-bottleneck trade-off and our observed trade-off, and reveal the over-compression phenomenon. Moreover, we investigate how and why the InfoNCE loss can alleviate the over-compression, and further present the contrastive temporal coding method. Our method successfully make discriminability and transferability compatible. Remarkable transfer learning performances are also achieved. We hope that this work can arouse attentions to the transferability of representations in the conventional supervised learning tasks. In the future, we will explore the existence of over-compression on other popular tasks, *e.g.*, self-supervised learning, object detection and large-scale pre-training.

Acknowledgements. This work was supported in part by the Zhejiang Provincial Natural Science Foundation of China under Grant No.LQ22F020006. We thank anonymous reviewers from ECCV 2022 for insightful comments.

References

- Alemi, A.A., Fischer, I., Dillon, J.V., Murphy, K.: Deep variational information bottleneck. ICLR (2017)
- Belghazi, M.I., Baratin, A., Rajeswar, S., Ozair, S., Bengio, Y., Courville, A., Hjelm, R.D.: MINE: mutual information neural estimation. arXiv:1801.04062 (2018)
- 3. Chen, C., Zheng, Z., Ding, X., Huang, Y., Dou, Q.: Harmonizing transferability and discriminability for adapting object detectors. In: CVPR (2020)
- 4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
- 5. Chen, T., Kornblith, S., Swersky, K., Norouzi, M., Hinton, G.: Big self-supervised models are strong semi-supervised learners. In: NeurIPS (2020)
- Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv:2003.04297 (2020)
- Chen, X., Wang, S., Long, M., Wang, J.: Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In: ICML (2019)
- Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: AISTATS (2011)
- 9. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: CVPR (2019)
- Darlow, L.N., Crowley, E.J., Antoniou, A., Storkey, A.J.: Cinic-10 is not imagenet or cifar-10. arXiv:1810.03505 (2018)
- 11. Feng, Y., Jiang, J., Tang, M., Jin, R., Gao, Y.: Rethinking supervised pre-training for better downstream transferring. In: ICLR (2022)
- 12. Fu, J., Zheng, H., Mei, T.: Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In: CVPR (2017)
- Furlanello, T., Lipton, Z., Tschannen, M., Itti, L., Anandkumar, A.: Born again neural networks. In: ICML (2018)
- 14. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
- 15. He, K., Girshick, R., Dollár, P.: Rethinking imagenet pre-training. In: ICCV (2019)
- 16. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv:1503.02531 (2015)
- Hjelm, R.D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., Bengio, Y.: Learning deep representations by mutual information estimation and maximization. In: ICLR (2019)
- 20. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: CVPR (2018)
- Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. ICLR (2020)
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. arXiv preprint arXiv:2004.11362 (2020)
- 23. Kornblith, S., Chen, T., Lee, H., Norouzi, M.: Why do better loss functions lead to less transferable features? In: NeurIPS (2021)
- 24. Kornblith, S., Shlens, J., Le, Q.V.: Do better imagenet models transfer better? In: CVPR (2019)

- 16 Q. Cui and B. Zhao et al.
- Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- 26. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: ECCV (2014)
- 28. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
- Long, M., Cao, Y., Wang, J., Jordan, M.: Learning transferable features with deep adaptation networks. In: ICML (2015)
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv:1306.5151 (2013)
- Mao, H., Chen, X., Fu, Q., Du, L., Han, S., Zhang, D.: Neuron campaign for initialization guided by information bottleneck theory. In: CIKM (2021)
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv:1807.03748 (2018)
- Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: ECCV (2020)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NeurIPS (2015)
- 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. IJCV (2015)
- Sariyildiz, M.B., Kalantidis, Y., Larlus, D., Alahari, K.: Concept generalization in visual representation learning. In: ICCV (2021)
- 37. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR (2015)
- Shao, J., Wen, X., Zhao, B., Xue, X.: Temporal context aggregation for video retrieval with contrastive learning. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2021)
- Shwartz-Ziv, R., Tishby, N.: Opening the black box of deep neural networks via information. arXiv:1703.00810 (2017)
- Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. arXiv:1906.05849 (2019)
- 41. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? In: NeurIPS (2020)
- Tishby, N., Zaslavsky, N.: Deep learning and the information bottleneck principle. In: ITW (2015)
- 43. Tripuraneni, N., Jordan, M., Jin, C.: On the theory of transfer learning: The importance of task diversity. In: NeurIPS (2020)
- 44. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: CVPR (2018)
- 45. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
- Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for self-supervised visual pre-training. arXiv:2011.09157 (2020)
- 47. Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: ECCV (2016)
- Wu, M., Zhuang, C., Mosse, M., Yamins, D., Goodman, N.: On mutual information in contrastive learning for visual representations. arXiv:2005.13149 (2020)

- 49. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via nonparametric instance discrimination. In: CVPR (2018)
- 50. Xiao, T., Li, S., Wang, B., Lin, L., Wang, X.: Joint detection and identification feature learning for person search. In: CVPR (2017)
- Yalniz, I.Z., Jégou, H., Chen, K., Paluri, M., Mahajan, D.: Billion-scale semisupervised learning for image classification. arXiv:1905.00546 (2019)
- 52. You, K., Liu, Y., Wang, J., Long, M.: Logme: Practical assessment of pre-trained models for transfer learning. In: ICML (2021)
- Zhao, B., Cui, Q., Song, R., Qiu, Y., Liang, J.: Decoupled knowledge distillation. In: CVPR (2022)
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017)
- 55. Zheng, H., Fu, J., Mei, T., Luo, J.: Learning multi-attention convolutional neural network for fine-grained image recognition. In: ICCV (2017)
- 56. Zhou, B., Cui, Q., Wei, X.S., Chen, Z.M.: BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In: CVPR (2020)
- 57. Zhu, R., Zhao, B., Liu, J., Sun, Z., Chen, C.W.: Improving contrastive learning by visualizing feature transformation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
- 58. Zoph, B., Ghiasi, G., Lin, T.Y., Cui, Y., Liu, H., Cubuk, E.D., Le, Q.V.: Rethinking pre-training and self-training. NeurIPS (2020)