

Understanding the Role of the Projector in Knowledge Distillation

Roy Miles, Krystian Mikołajczyk

Imperial College London
r.miles18@imperial.ac.uk, k.mikolajczyk@imperial.ac.uk

Abstract

In this paper we revisit the efficacy of knowledge distillation as a function matching and metric learning problem. In doing so we verify three important design decisions, namely the normalisation, soft maximum function, and projection layers as key ingredients. We theoretically show that the projector implicitly encodes information on past examples, enabling relational gradients for the student. We then show that the normalisation of representations is tightly coupled with the training dynamics of this projector, which can have a large impact on the students performance. Finally, we show that a simple soft maximum function can be used to address any significant capacity gap problems. Experimental results on various benchmark datasets demonstrate that using these insights can lead to superior or comparable performance to state-of-the-art knowledge distillation techniques, despite being much more computationally efficient. In particular, we obtain these results across image classification (CIFAR100 and ImageNet), object detection (COCO2017), and on more difficult distillation objectives, such as training data efficient transformers, whereby we attain a 77.2% top-1 accuracy with DeiT-Ti on ImageNet. Code and models are publicly available.

Introduction

Deep neural networks have achieved remarkable success in various applications, ranging from computer vision (Krizhevsky, Sutskever, and Hinton 2012) to natural language processing (Vaswani et al. 2017). However, the high computational cost and memory requirements of deep models have limited their deployment in resource-constrained environments. Knowledge distillation is a popular technique to address this problem through transferring the knowledge of a large teacher model to that of a smaller student model. This technique involves training the student to imitate the output of the teacher, either by directly minimizing the difference between intermediate features or by minimizing the Kullback-Leibler (KL) divergence between their soft predictions. Although knowledge distillation has shown to be very effective, there are still some limitations related to the computational and memory overheads in constructing and evaluating the losses, as well as an insufficient theoretical explanation for the underlying core principles.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

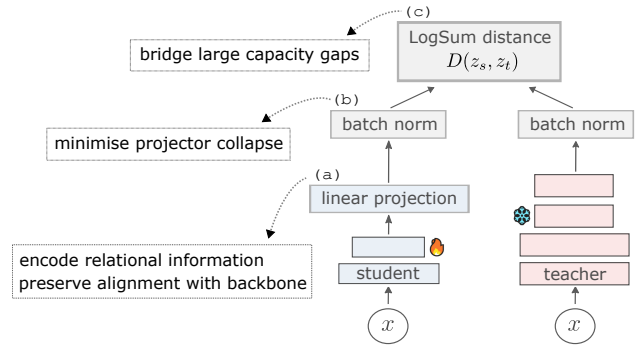


Figure 1: Proposed feature distillation pipeline using three distinct components: linear projection (a), batch norm (b), and a *LogSum* distance (c). We provide an interpretable explanation for each each of these three components, which results in a very cheap and effective recipe for distillation.

To overcome these limitations, we revisit knowledge distillation from both a function matching and metric learning perspective. We perform an extensive ablation of three important components of knowledge distillation, namely the distance metric, normalisation, and projector network. Alongside this ablation we provide a theoretical perspective and unification of these design principles through exploring the underlying training dynamics. Finally, we extend these principles to a few large scale vision tasks, whereby we achieve comparable or improved performance over state-of-the-art. The most significant result of which pertains to the data-efficient training of transformers, whereby a performance gain of 2.2% is achieved over the best-performing distillation methods that are designed explicitly for this task. Our main contributions can be summarised as follows.

- We explore three distinct design principles from knowledge distillation, namely the projection, normalisation, and distance function. In doing so we demonstrate their and coupling with each other, both through analytical means and by observing the training dynamics.
- We show that a projection layer implicitly encodes relational information from previous samples. Using this knowledge we can remove the need to explicitly construct correlation matrices or memory banks that will inevitably incur a significant memory overhead.

- We propose a simple recipe for knowledge distillation using a linear projection, batch normalisation, and a *LogSum* function. These three design choices can attain competitive or improved performance to state-of-the-art for image classification, object detection, and the data efficient training of transformers.

Related Work

Knowledge Distillation Knowledge distillation is the process of transferring the knowledge from a large, complex model to a smaller, simpler model. Its usage was originally proposed in the context of image classification (Hinton, Vinyals, and Dean 2015) whereby the soft teacher predictions would encode relational information between classes. Spherical KD (Guo et al. 2020) extended this idea by re-scaling the logits, prime aware adaptive distillation (Zhang et al. 2020) introduced an adaptive weighting strategy, while DKD (Zhao, Song, and Qiu 2022) proposed to decouple the original formulation into target class and non-target class probabilities.

Hinted losses (Romero et al. 2015) were a natural extension of the logit-based approach whereby the intermediate feature maps are used as hints for the student. Attention transfer (Zagoruyko and Komodakis 2019) then proposed to re-weight this loss using spatial attention maps. ReviewKD (Chen et al. 2021a) addressed the problem relating to the arbitrary selection of layers by aggregating information across all the layers using trainable attention blocks. Neuron selectivity transfer (Huang and Wang 2017), similarity-preserving KD (Tung and Mori 2019), and relational KD (Park et al. 2019) construct relational batch and feature matrices that can be used as inputs for the distillation losses. Similarly FSP matrices (Yim 2017) were proposed to extract the relational information through a residual block. In contrast to this theme, we show that a simple projection layer can implicitly capture most relational information, thus removing the need to construct any expensive relational structures.

Representation distillation was originally proposed alongside a contrastive based loss (Tian, Krishnan, and Isola 2019) and has since been extended using a Wasserstein distance (Chen et al. 2020a), information theory (Miles, Rodriguez, and Mikolajczyk 2022), graph theory (Ma, Chen, and Akata 2022), and complementary gradient information (Zhu et al. 2021a). Distillation also been empirically shown to benefit from longer training schedules and more data-augmentation (Beyer et al. 2022), which is similarly observed with HSAKD (Yang et al. 2021) and SSKD (Xu et al. 2020). Distillation between CNNs and transformers has also been a very practically motivated task for data-efficient training (Touvron et al. 2021) and has shown to benefit from an ensemble of teacher architectures (Ren et al. 2022). However, we show that just a simple extension of some fundamental distillation design principles is much more effective.

Self-distillation is another branch of knowledge distillation that instead proposes to instead distill knowledge within the network itself (Zhang et al. 2019). This paradigm removes the need to have any pre-trained teacher readily available and has since been applied in the context of deep metric

learning (Roth et al. 2021), graph neural networks (Chen et al. 2021b), and image classification (Miles and Mikolajczyk 2020). Its success of which has since driven many theoretical advancements (Mobahi, Farajtabar, and Bartlett 2020; Allen-Zhu and Li 2023; Zhang and Sabuncu 2020) and is still an active area of research.

Self-Supervised Learning Self-supervised learning (SSL) is an increasingly popular field of machine learning whereby a model is trained to learn a useful representation of unlabelled data. Its popularity has been driven by the increasing cost of manual labelling and has since been crucial for training large transformer models. Various pretext tasks have been proposed to learn these representations, such as image inpainting (He et al. 2022), colorization (Zhang, Isola, and Efros 2016), or prediction of the rotation (Gidaris, Singh, and Komodakis 2018) or position of patches (Doersch, Gupta, and Efros 2015; Carlucci et al. 2019). SimCLR (Chen et al. 2020b) approached self-supervision using a contrastive loss with multi-view augmentation to define the positive and negative pairs. They found a large memory bank of negative representations was necessary to achieve good performance, but would incur a significant memory overhead. MoCo (He et al. 2020) extending this work with a momentum encoder, which was subsequently extended by MoCov2 (Chen et al. 2020c) and MoCov3 (Chen, Xie, and He 2021).

Feature decorrelation is another approach to SSL that avoids the need for negative pairs to address representation collapse. Both Barlow twins (Zbontar et al. 2021) and VICReg (Bardes, Ponce, and LeCun 2022a) achieve this by maximising the variance within a batch, while preserving invariance to augmentations. Both contrastive learning and feature decorrelation have since been extended to dense prediction tasks (Bardes, Ponce, and LeCun 2022b) and unified with knowledge distillation (Miles et al. 2023).

Understanding the Role of the Projector

Knowledge distillation (KD) is a technique used to transfer knowledge from a large, powerful model (teacher) to a smaller, less powerful one (student). In the classification setting, it can be done using the soft teacher predictions as pseudo-labels for the student. Unfortunately, this approach does not trivially generalise to non-classification tasks (Liu et al. 2019) and the classifier may collapse a lot of information (Tishby 2015) that can be useful for distillation. Another approach is to use feature maps from the earlier layers for distillation (Romero et al. 2015; Zagoruyko and Komodakis 2019), however, its usage presents two primary challenges: the difficulty in ensuring consistency across different architectures (Chen et al. 2021a) and the potential degradation in the student’s downstream performance for cases where the inductive biases of the two networks differ (Tian, Krishnan, and Isola 2019). A compromise, which strikes a balance between the two approaches discussed above, is to distill the representation directly before the output space. Representation distillation has been successfully adopted in past works (Tian, Krishnan, and Isola 2019; Zhu et al. 2021a; Miles, Rodriguez, and Mikolajczyk 2022) and is the focus of this paper. The exact training framework used is described in figure 1. The




Figure 2: Evolution of singular values of the projection weights \mathbf{W}_p under three different representation normalisation schemes. The student is a Resnet-18, while the teacher is a ResNet-50. The three curves shows the evolution of singular values for the projector weights when the representations undergo no normalisation, L2 normalisation, and batch norm respectively.

projection layer shown was originally used to simply match the student and teacher dimensions (Romero et al. 2015), however, we will show that its role is much more important and it can lead to significant performance improvements even when the two feature dimensions already match. The two representations are typically both followed by some normalisation scheme as a way of appropriately scaling the gradients. However, we find this normalisation has a more interesting property in its relation to what information is encoded in the learned projector weights.

In this work we provide a theoretical perspective to motivate some simple and effective design choices for knowledge distillation. In contrast to the recent works (Tung and Mori 2019; Miles, Rodriguez, and Mikolajczyk 2022), we show that an explicit construction of complex relational structures, such as feature kernels (He and Ozay 2022) is not necessary. In fact, most of this structure can be learned implicitly through the tightly coupled interaction of a learnable projection layer and an appropriate normalisation scheme. In the following sections we investigate the training dynamics of the projection layer with the choice of normalisation scheme. We explore the impact and trade-offs that arise from the architecture design of the projector. Finally, we propose a simple modification to the distance metric to address issues arising from a large capacity gap between the student and teacher models. Although we do not aim to necessarily propose a new method for distillation, we uncover a cheap and simple recipe that can transfer to various distillation settings and tasks. Furthermore, we provide a new theoretical perspective on the underlying principles of distillation that can translate to large scale vision tasks.

The projection weights encode relational information from previous samples. The projection layer plays a crucial role in KD as it provides an implicit encoding of previous samples and its weights can capture the relational information needed to transfer information regarding the correlation between features. We observe that even a single linear projector layer can provide significant improvements in accuracy (see Supplementary). This improvement suggests that the pro-

jections role in distillation can be described more concisely as being an encoder of essential information needed for the distillation loss itself. Most recent works propose a manual construction of some relational information to be used as part of a loss (Park et al. 2019; Tung and Mori 2019), however, we posit that an implicit and learnable approach is much more effective. To explore this phenomenon in more detail, we consider the update equations for the projector weights and its training dynamics. Without loss in generality, consider a simple L2 loss and a linear bias-free projection layer.

$$D(\mathbf{Z}_s, \mathbf{Z}_t; \mathbf{W}_p) = \frac{1}{2} \|\mathbf{Z}_s \mathbf{W}_p - \mathbf{Z}_t\|_2^2 \quad (1)$$

Where \mathbf{Z}_s and \mathbf{Z}_t are the student and teacher representations respectively, while \mathbf{W}_p is the matrix representing the linear projection. Using the trace property of the Frobenius norm, we can then express this loss as follows:

$$D(\mathbf{Z}_s, \mathbf{Z}_t; \mathbf{W}_p) = \frac{1}{2} \text{tr} \left((\mathbf{Z}_s \mathbf{W}_p - \mathbf{Z}_t)^T (\mathbf{Z}_s \mathbf{W}_p - \mathbf{Z}_t) \right) \quad (2)$$

$$= \frac{1}{2} \text{tr} (\mathbf{W}_p^T \mathbf{Z}_s^T \mathbf{Z}_s \mathbf{W}_p - \mathbf{Z}_t^T \mathbf{Z}_s \mathbf{W}_p \quad (3)$$

$$- \mathbf{W}_p^T \mathbf{Z}_s^T \mathbf{Z}_t + \mathbf{Z}_t^T \mathbf{Z}_t) \quad (4)$$

Taking the derivative with respect to \mathbf{W}_p , we can derive the update rule $\dot{\mathbf{W}}_p$

$$\dot{\mathbf{W}}_p = -\frac{\partial D(\mathbf{W}_p)}{\partial \mathbf{W}_p} = -\mathbf{Z}_s^T \mathbf{Z}_s \mathbf{W}_p + \mathbf{Z}_s^T \mathbf{Z}_t \quad (5)$$

which can be further simplified

$$\dot{\mathbf{W}}_p = \mathbf{C}_{st} - \mathbf{C}_s \mathbf{W}_p \quad (6)$$

where $\mathbf{C}_s = \mathbf{Z}_s^T \mathbf{Z}_s \in \mathbb{R}^{d_s \times d_s}$ and $\mathbf{C}_{st} = \mathbf{Z}_s^T \mathbf{Z}_t \in \mathbb{R}^{d_s \times d_t}$ denote self and cross correlation matrices respectively. Due to the capacity gap between the student network and the teacher network, there is no perfect linear projection between these two representation spaces. Instead, the projector will

converge on an approximate projection that we later show is governed by the normalisation being employed.

Whitened features: consider using self-supervised learning in conjunction with distillation whereby the student features are whitened to have perfect decorrelation (Ermolov et al. 2020), or alternatively, they are batch normalised and sufficiently regularised with a feature decorrelation term (Bardes, Ponce, and LeCun 2022a). In this setting, the fixed point solution for the weights will be symmetric and will capture the cross relationship between student and teacher features.

$$\mathbf{C}_{st} - \mathbf{C}_s \mathbf{W}_p = 0 \quad \text{where} \quad \mathbf{C}_s = \mathbf{I} \quad (7)$$

$$\rightarrow \mathbf{W}_p = \mathbf{C}_{st} \quad (8)$$

Other normalisation schemes, such as those that jointly normalise the projected features and the teacher features, will have a much more involved analysis but will unlikely provide any additional insights on the dynamics of training itself. Thus, we propose to empirically explore the training trajectories of the projector weights singular values. This exploration will help quantify how the projector is mapping the student features to the teachers space. We cover this in the next section along with some additional insights into what is being learned and distilled.

The choice of normalisation directly affects the training dynamics of \mathbf{W}_p . Equation 6 shows that the projector weights can encode relational information between the student and teacher’s features. This suggests redundancy in explicitly constructing and updating a large memory bank of previous representations (Tian, Krishnan, and Isola 2019). By considering a weight decay η and a learning rate α_p , the update equation can be given as follows:

$$\mathbf{W}_p \rightarrow \mathbf{W}_p + \alpha_p \dot{\mathbf{W}}_p - \eta \mathbf{W}_p \quad (9)$$

$$= (1 - \eta) \mathbf{W}_p + \alpha_p \dot{\mathbf{W}}_p \quad (10)$$

By setting $\eta = \alpha_p$ we can see that the projection layer will reduce to a moving average of relational features, which is very similar to the momentum encoder used by CRD (Tian, Krishnan, and Isola 2019). Other works suggest to extract relational information on-the-fly by constructing correlation or Gram matrices (Miles, Rodriguez, and Mikolajczyk 2022; Peng et al. 2019). We show that this is also not necessary and more complex information can be captured through a simple linear projector. We also demonstrate that, in general, the use of a projector will scale much more favourably for larger batch sizes and feature dimensions. We also note that the handcrafted design of kernel functions (Joshi et al. 2021; He and Ozay 2022) may not generalise to real large scale datasets without significant hyperparameter tuning.

From the results in table 1, we observe that when fixing all other settings, the choice of normalisation can significantly affect the student’s performance. To explore this in more detail, we consider the training trajectories of \mathbf{W}_p under different normalisation schemes. We find that the choice of normalisation not only controls the training dynamics, but also the fixed point solution (see equation 6). We argue that the efficacy of distillation is dependent on how much relational information can be encoded in the learned weights

and how much information is lost through the projection. To jointly evaluate these two properties we show the evolution of singular values of the projector weights during training. The results can be seen in figure 2 and show that the better performing normalisation methods (table 1) are shrinking far fewer singular values towards zero. This shrinkage can be described as collapsing the input along some dimension, which will induce some information loss and it is this information loss that degenerates the efficacy of the distillation process.

Teacher	RegNet	ViT	ConvNeXt
Student	MBv2	MBv3	EffNet-b0
No distillation	50.89	54.13	64.48
L2 Norm	52.91	54.65	64.23
Group Norm	55.63	59.08	65.52
Batch Norm	56.09	59.28	67.95

Table 1: Normalisation ablation for distillation across a range of architecture pairs on ImageNet-1K 20% subset.

Larger projector networks learn to decorrelate the input-output features. One natural extension of the previous observations is to use a larger projector network to encode more information relevant for the distillation loss. Unfortunately, we observe that a trivial expansion of the projection architecture does not necessarily improve the students performance. To explain this observation we evaluate a measure of decorrelation between the input and output features of these projector networks. The results can be seen in figure 3 and we can see that the larger projectors learn to decorrelate more and more features from the input. This decorrelation can lead to the projector learning features that are not shared with the student backbone, which will subsequently diminish the effectiveness of distillation. These observations suggest that there is an inherent trade-off between the projector capacity and the efficacy of distillation. We note that the handcrafted design of the projector architecture is a motivated direction for further research (Chen et al. 2022b; Navaneet et al. 2021). However, in favour of simplicity, we choose to use a linear projector for all of our large scale evaluations.




Figure 3: Correlation between input-output features using different projector architectures. All projectors considered will gradually decorrelate the input-output features.

Teacher	ViT	ConvNeXt	ResNet50
Student	MBv3	EffNet-b0	ResNet18
w/ LogSum	59.28	67.95	70.03
w/ LogSum	59.80	68.51	71.29

Table 2: LogSum ablation across various architecture pairs. Left: 20% subset. Right: Full ImageNet. The soft maximum function provides consistent improvement across both the CNN→CNN and ViT→CNN distillation settings.

The soft maximum function can address distilling across a large capacity gap. When the capacity gap between the student and the teacher is large, representation distillation can become challenging. More specifically, the student network may have insufficient capacity to perfectly align these two spaces and in attempting to do so may degrade its downstream performance. To address this issue we explore the use of a soft maximum function which will soften the contribution of relatively close matches in a batch. In this way the loss can be adjusted to compensate for poorly aligned features which may arise when the capacity gap is large. The family of functions which share these properties can be more broadly defined through a property of their gradients. In favour of simplicity, we use the simple *LogSum* function throughout our experiments.

$$D(\mathbf{Z}_s, \mathbf{Z}_t; \mathbf{W}_p) = \log \sum_i |\mathbf{Z}_s \mathbf{W}_p - \mathbf{Z}_t|_i^\alpha \quad (11)$$

where α is a smoothing factor. We also note that other functions, such as the *LogSumExp*, with a temperature parameter τ , have been used in SimCLR and CRD to a similar effect. Table 1 shows the importance of feature normalisation across a variety of student-teacher architecture pairs. Batch normalisation provides the most consistent improvement that even extends to the Transformer → CNN setting. In table 2 we highlight the importance of the *LogSum* function, which is most effective in the large capacity gap settings, as evident from the 1% improvement for R50 → R18. Table 3 provides an ablation of the importance of the α parameter, whereby we observe that the performance is relatively robust to a wide range of values, but consistently optimal in the range 4-5.

Teacher	ResNet50	ConvNeXt
Student	ResNet18	EffNet-b0
1.0	61.74	65.52
2.0	62.23	66.61
3.0	63.10	67.72
4.0	63.32	68.51
5.0	63.40	67.69

Table 3: Ablating the importance of α . Distillation is generally robust for various values of α , but consistently optimal in range 4-5 across various architecture pairs.

Benchmark Evaluation

Implementation details. We follow the same training schedule as CRD (Tian, Krishnan, and Isola 2019) for both the CIFAR100 and ImageNet experiments. For the object detection, we use the same training schedule as ReviewKD (Chen et al. 2021a), while for the data efficient training we use the same as Co-Advice (Ren et al. 2022). All experiments were performed on a single NVIDIA RTX A5000. When using batch normalisation for the representations, we removed the affine parameters and set $\epsilon = 0.0001$. For all experiments we jointly train the student using a task loss \mathcal{L}_{task} and the feature distillation loss given in equation 11.

$$\mathcal{L} = \mathcal{L}_{task} + D(\mathbf{Z}_s, \mathbf{Z}_t; \mathbf{W}_p) \quad (12)$$

Data Efficient Training for Transformers

Transformers have emerged as a viable replacement for convolution-based neural networks (CNN) in visual learning tasks. Despite the promise of these models, their performance will suffer when there is insufficient training data available, such as in the case of ImageNet. DeiT (Touvron et al. 2021) was the first to address this problem through the use of knowledge distillation. Although the authors show improved alignment with the teacher, we believe this fails to capture why less data is needed.

Network	acc@1	Teacher	#params
RegNetY-160	82.6	none	84M
BiT-M R152x2	84.5	none	236M
DeiT-Ti	72.2	none	5M
CivT-Ti	74.9	ensemble	6M
DeiT-Ti	74.5	regnety-160	6M
↳ 1000 epochs	76.6	regnety-160	6M
DearKD	74.8	regnety-160	6M
↳ 1000 epochs	77.0	regnety-160	6M
USKD	75.0	regnety-160	6M
Our Method	77.2	regnety-160	6M
ResNet-50	76.5	none	25M
FunMatch	80.3	bit-m r152x2	25M
↳ 9600 epochs	82.8	bit-m r152x2	25M
DeiT-S	79.8	none	22M
CivT-S	82.0	ensemble	22M
DeiT-S	81.2	regnety-160	22M
↳ 1000 epochs	82.6	regnety-160	22M
DearKD	81.5	regnety-160	22M
↳ 1000 epochs	82.8	regnety-160	22M
USKD	80.8	regnety-160	22M
Our Method	82.1	regnety-160	22M

Table 4: Data-efficient training of transformers and CNNs on the ImageNet-1K dataset. Unless specified, all student models are trained for 300 epochs.

We posit that the distillation process encourages the student to learn layers which are "more" translational equivariant in attempt to match the teacher's underlying function. Although this is the principle that motivates using an ensemble of teacher models with different inductive biases (Ren

Teacher	WRN40-2	WRN40-2	R56	R32x4	VGG13	R50	R50	R32x4	R32x4	WRN40-2
Student	WRN16-2	WRN40-1	R20	R8x4	MBv2	MBv2	VGG8	ShuffleV1	ShuffleV2	ShuffleV1
Teacher	76.46	76.46	73.44	79.63	75.38	79.10	79.10	79.63	79.63	76.46
Student	73.64	72.24	69.63	72.51	65.79	65.79	70.68	70.77	73.12	70.77
KD	74.92	73.54	70.66	73.33	67.37	67.35	73.81	74.07	74.45	74.83
FitNet	75.75	74.12	71.60	74.31	68.58	68.54	73.84	74.82	75.11	75.55
AT	75.28	74.45	71.78	74.26	69.34	69.28	73.45	74.76	75.30	75.61
CRD	76.04	75.52	71.68	75.90	68.49	70.32	74.42	75.46	75.72	75.96
SSKD	76.04	76.13	71.49	76.20	71.53	72.57	75.76	78.44	78.61	77.40
Our Method	76.14	75.42	71.75	76.44	71.47	72.81	76.20	77.32	79.06	79.22
KD [†]	75.94	75.32	71.10	75.84	70.79	71.29	75.75	77.80	78.43	78.00
CRD [†]	77.27	76.15	72.21	77.69	71.65	72.03	75.73	78.57	79.01	78.54
DKD [†]	74.96	75.89	70.95	77.52	72.01	73.30	76.88	79.71	80.08	77.86
Our Method	77.61	76.04	72.25	78.37	72.82	73.51	77.08	78.99	79.86	78.79

Table 5: KD between Similar and Different Architectures. Top-1 accuracy (%) on CIFAR100. Bold is used to denote the best results. All reported models are trained using pairs of augmented images. Those reported in the top box use RandAugment (Cubuk et al. 2020) strategy, while those in the bottom box use pre-defined rotations, as used in SSKD. [†] denotes reproduced results in a new augmentation setting using the authors provided code.

et al. 2022), there is still no thorough demonstration on if the inductive biases are actually being transferred. In this section we attempt to address this gap by introducing a measure of equivariance. We show that applying our distillation principles to this task can achieve significant improvements over state-of-the-art as a result of transferring more of the translational equivariance.

The results of these experiments are shown in table 4. We use the exact same training methodology as co-advice (Ren et al. 2022) and choose to use batch normalisation, a linear projection layer, and $\alpha = 4$ as the parameters for distillation. We observe a significant improvement over both DeiT and CivT when the capacity gap is large. However, as the capacity gap diminishes, and the student approaches the same performance as the teacher, this improvement is much less significant. Multiple factors, such as the soft maximum function and the batch normalisation, will be contributing to this observed result. However, the explanation is more concisely described by the fact that our distillation loss transfers more translational equivariance to the student.

Classification on CIFAR100 and ImageNet

Experiments on the CIFAR-100 classification task (Krizhevsky 2009) consist of 60K 32x32 RGB images across 100 classes with a 5:1 training/testing split. Table 5 shows the results for several student-teacher pairings. To enable a fair evaluation, we have only included the methods that use the same teacher weights provided by SSKD (Xu et al. 2020). In these experiments we use an MLP projector with a hidden size of 1024 and no additional KL divergence loss. We confirm that not only is the choice of augmentation critical for good performance (Beyer et al. 2022) on this dataset, but applying our principles can attain state-of-the-art across most architecture pairs. The most significant improvements pertain to the cross-architecture experiments or where the capacity gap is large. We provide two sets of experiments with and without introducing a wider

set of augmentations. In both settings we maintain the same optimiser, learning rate scheduler, and training duration.

The ImageNet (Russakovsky et al. 2014) classification uses 1.3 million images that are classified into 1000 distinct classes. The input size are set to 224 x 224, and we employed a typical augmentation procedure that includes cropping and horizontal flipping. We used the torchdistill library with the standard configuration, which involves 100 training epochs using SGD and an initial learning rate of 0.1, which is decreased by a factor of 10 at epochs 30, 60, and 90. The results can be seen in table 6 and although the choice of architectures is not in favour of our method since the capacity gap is small, we are still able to attain competitive performance. Other methods, such as ICKD (Liu et al. 2021) or SimKD (Chen et al. 2022a) either modify the original training settings or architectures, and so have been omitted from this evaluation.

Cross architecture distillation can implicitly transfer inductive biases. CNNs use convolutions, which are spatially local operations, whereas transformers use self-attention, which are global operations. We expect that a benefit of this cross-architecture distillation setting is that the students learn to be "more" spatially equivariant in an attempt to match the teachers underlying function. It is this strong inductive bias that can reduce the amount of training data needed. A layer is translation equivariant if the following property holds:

$$\phi(T\mathbf{x}) = T\phi(\mathbf{x}) \quad (13)$$

In other words, if we take a translated input $T\mathbf{x}$ and pass it through a layer ϕ , the result should be equivalent to first applying ϕ to \mathbf{x} and then performing the translation. A natural measure of equivariance can then be the difference between the left and right-hand side of this equation 13.

$$\mu_T(\phi) = \|\phi(T\mathbf{x}) - T\phi(\mathbf{x})\|_2^2 \quad (14)$$

	Teacher	Student	AT	KD	CC	CRD	ReviewKD	Ours
acc@1	26.69	30.25	29.30	29.34	30.04	28.62	28.39	28.37
acc@5	8.58	10.93	10.00	10.12	10.83	9.51	9.42	9.41

Table 6: Top-1 and Top-5 error rates (%) on ImageNet. ResNet18 as student, ResNet34 as teacher.

We evaluate this measure on a block of self-attention layers by first removing the distillation and class tokens and then rolling the patch tokens to recover the spatial dimensions. This operation can then be performed on the input and output tensors before applying a translation. Table 7 shows this measure of equivariance after training with and without distillation. In general, we observe that the distilled models do in fact learn to preserve spatial locality between feature maps, which aligns with the function matching perspective for distillation.

We find that our simple distillation recipe can transfer a lot more of this equivariance property to the student. Although Co-advise does learn this spatial locality to some extent, it is much less significant than using our feature based distillation, despite both attaining a similar level of performance. Intermediate feature map losses may be able to transfer even more of this translational equivariance property, however, its usage may degrade the benefit of using a transformer in the first place. For example, although we observe that most self-attention blocks (trained using distillation) do preserve a lot of this spatial locality, there is still some global context between patch tokens that is still being preserved.

Network	$\mu_T(\phi)$
DeiT-S	1.52 ± 0.15
CivT-S	0.13 ± 0.05
Our Method	0.04 ± 0.02

Table 7: Measure of translational equivariance for a DeiT-S transformer model trained with and without distillation. These results confirm that distillation can transfer explicit inductive biases from the teacher.

Object Detection on COCO

We extend the application of our method to object detection, whereby we employ a similar approach as used in the classification task by distilling the backbone output features of both the student and teacher networks. Due to the smaller batch sizes used in these experiments, we choose to instead normalise over the height and width dimensions. For evaluating the efficacy of our method, we conduct experiments on the widely-used COCO2017 dataset (Lin et al. 2014) under the same settings provided in ReviewKD. We then further demonstrate the applicability of our distillation principles on the more recent and efficient YOLOv5 model (Zhu et al. 2021b). In both cases we show improved student performance on the downstream task, whereby competitive performance is achieved with ReviewKD despite being significantly simpler and cheaper to integrate into a given distillation pipeline. Our method even outperforms FPGI (Wang et al. 2019), which is directly designed for detection.

Model	mAP (50-95)	mAP 50
YOLOv5m (teacher)	64.1	45.4
YOLOv5s	56.8	37.4
↳ Our Method	57.3	37.5
	mAP	AP50
Faster R-CNN w/ R50 (teacher)	40.22	61.02
Faster R-CNN w/ MV2	29.47	48.87
↳ KD	30.13	50.28
↳ FitNet	30.20	49.80
↳ FPGI	31.16	50.68
↳ ReviewKD	33.71	53.15
↳ Our Method	<u>32.92</u>	<u>52.96</u>

Table 8: Object detection on COCO. (top) We report the standard COCO metric of mAP averaged over IOU thresholds in [0.5 : 0.05 : 0.95] along with the standard PASCAL VOC’s metric (Everingham et al. 2010), which is the average mAP@0.5. (bottom) For the R-CNN results, we report the mAP and AP50 metrics to enable a consistent comparison with ReviewKD.

Conclusion

In this paper, we revisited the core underlying principles of knowledge distillation and have performed an extensive ablation on the most effective and scaleable components. In doing so, we have provided a new theoretical perspective for understanding these results through analyzing the projector training dynamics. By extending these principles to a wide range of tasks, we achieve competitive or improved performance to state-of-the-art across image classification, object detection, and data efficient training of transformers. Our proposed distillation recipe can significantly reduce the complexity and memory consumption of existing pipelines by avoiding the need to construct expensive relational object, many trainable layers, or enforcing very long training schedules. We further show improved performance for the large capacity gap settings and evidence for the distillation of explicit inductive biases from the teacher. Looking ahead to future research in this area, we expect to see the joint development of more sophisticated normalisation schemes and projection networks, which will encode more complex and informative features for the distillation process.

Code Reproducibility. To facilitate the reproducibility of results, we release all the training code and pre-trained weights. The ImageNet experiments are performed using the popular *torchdistill* (Matsubara 2020) framework, while the CIFAR100 and data-efficient training code are based on those provided by CRD (Tian, Krishnan, and Isola 2019) and co-advise (Ren et al. 2022) respectively.

References

- Allen-Zhu, Z.; and Li, Y. 2023. Towards Understanding Ensemble, Knowledge Distillation and Self-Distillation in Deep Learning.
- Bardes, A.; Ponce, J.; and LeCun, Y. 2022a. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. *ICLR*.
- Bardes, A.; Ponce, J.; and LeCun, Y. 2022b. VICRegL: Self-Supervised Learning of Local Visual Features. *NeurIPS*.
- Beyer, L.; Zhai, X.; Royer, A.; Markeeva, L.; Anil, R.; and Kolesnikov, A. 2022. Knowledge distillation: A good teacher is patient and consistent. *CVPR*.
- Carlucci, F. M.; D’Innocente, A.; Bucci, S.; Caputo, B.; and Tommasi, T. 2019. Domain Generalization by Solving Jigsaw Puzzles. *CVPR*.
- Chen, D.; Mei, J.-P.; Zhang, H.; Wang, C.; Feng, Y.; and Chen, C. 2022a. Knowledge Distillation with the Reused Teacher Classifier. *CVPR*.
- Chen, L.; Wang, D.; Gan, Z.; Liu, J.; Henao, R.; and Carin, L. 2020a. Wasserstein Contrastive Representation Distillation. *CVPR*.
- Chen, P.; Liu, S.; Zhao, H.; and Jia, J. 2021a. Distilling Knowledge via Knowledge Review. *CVPR*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. *ICML*.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020c. Improved Baselines with Momentum Contrastive Learning. *arXiv preprint*.
- Chen, X.; Xie, S.; and He, K. 2021. An Empirical Study of Training Self-Supervised Vision Transformers. *ICCV*.
- Chen, Y.; Bian, Y.; Xiao, X.; Rong, Y.; Xu, T.; and Huang, J. 2021b. On Self-Distilling Graph Neural Network.
- Chen, Y.; Wang, S.; Liu, J.; Xu, X.; de Hoog, F.; and Huang, Z. 2022b. Improved Feature Distillation via Projector Ensemble. *NeurIPS*.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. *CVPR Workshop*.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised Visual Representation Learning by Context Prediction. *ICCV*.
- Ermolov, A.; Siarohin, A.; Sangineto, E.; and Sebe, N. 2020. Whitening for Self-Supervised Representation Learning. *ICML*.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised Representation Learning by Predicting Image Rotations. *ICLR*.
- Guo, J.; Chen, M.; Hu, Y.; Zhu, C.; He, X.; and Cai, D. 2020. Reducing the Teacher-Student Gap via Spherical Knowledge Distillation. *arXiv preprint*.
- He, B.; and Ozay, M. 2022. Feature Kernel Distillation. *ICLR*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked Autoencoders Are Scalable Vision Learners. *CVPR*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum Contrast for Unsupervised Visual Representation Learning. *CVPR*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *NeurIPS*.
- Huang, Z.; and Wang, N. 2017. Like What You Like: Knowledge Distill via Neuron Selectivity Transfer. *arXiv preprint*.
- Joshi, C. K.; Liu, F.; Xun, X.; Lin, J.; and Foo, C.-S. 2021. On Representation Knowledge Distillation for Graph Neural Networks. *arXiv preprint*.
- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. *NeurIPS*.
- Lin, T. Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common objects in context. *ECCV*.
- Liu, L.; Huang, Q.; Lin, S.; Xie, H.; Wang, B.; Chang, X.; and Liang, X. 2021. Exploring Inter-Channel Correlation for Diversity-preserved Knowledge Distillation. *ICCV*.
- Liu, Y.; Chen, K.; Liu, C.; Qin, Z.; Luo, Z.; and Wang, J. 2019. Structured Knowledge Distillation for Semantic Segmentation. *CVPR*.
- Ma, Y.; Chen, Y.; and Akata, Z. 2022. Distilling Knowledge from Self-Supervised Teacher by Embedding Graph Alignment. *BMVC*.
- Matsubara, Y. 2020. torchdistill : A Modular, Configuration-Driven Framework for Knowledge Distillation.
- Miles, R.; and Mikolajczyk, K. 2020. Cascaded channel pruning using hierarchical self-distillation. *BMVC*.
- Miles, R.; Rodriguez, A. L.; and Mikolajczyk, K. 2022. Information Theoretic Representation Distillation. *BMVC*.
- Miles, R.; Yucel, M. K.; Manganelli, B.; and Saa-Garriga, A. 2023. MobileVOS: Real-Time Video Object Segmentation Contrastive Learning meets Knowledge Distillation. *CVPR*.
- Mobahi, H.; Farajtabar, M.; and Bartlett, P. L. 2020. Self-Distillation Amplifies Regularization in Hilbert Space.
- Navaneet, K. L.; Koohpayegani, S. A.; Tejankar, A.; and Pirsiavash, H. 2021. SimReg: Regression as a Simple Yet Effective Tool for Self-supervised Knowledge Distillation. *BMVC*.
- Park, W.; Corp, K.; Kim, D.; and Lu, Y. 2019. Relational Knowledge Distillation. *CVPR*.
- Peng, B.; Jin, X.; Li, D.; Zhou, S.; Wu, Y.; Liu, J.; Zhang, Z.; and Liu, Y. 2019. Correlation congruence for knowledge distillation. *CVPR*.
- Ren, S.; Gao, Z.; Hua, T.; Xue, Z.; Tian, Y.; He, S.; and Zhao, H. 2022. Co-advise: Cross Inductive Bias Distillation. *CVPR*.
- Romero, A.; Ballas, N.; Ebrahimi Kahou, S.; Chassang, A.; Gatta, C.; and Bengio, Y. 2015. FitNets: Hints For Thin Deep Nets. *ICLR*.

- Roth, K.; Milbich, T.; Ommer, B.; Cohen, J. P.; and Ghassemi, M. 2021. S2SD: Simultaneous Similarity-based Self-Distillation for Deep Metric Learning.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2014. ImageNet Large Scale Visual Recognition Challenge. *IJCV*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive representation distillation. *ICLR*.
- Tishby, N. 2015. Deep Learning and the Information Bottleneck Principle. *IEEE Information Theory Workshop (ITW)*.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers & distillation through attention. *PMLR*.
- Tung, F.; and Mori, G. 2019. Similarity-preserving knowledge distillation. *ICCV*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. *NeurIPS*.
- Wang, T.; Yuan, L.; Zhang, X.; and Feng, J. 2019. Distilling Object Detectors with Fine-grained Feature Imitation. *CVPR*.
- Xu, G.; Liu, Z.; Li, X.; and Loy, C. C. 2020. Knowledge Distillation Meets Self-supervision. *ECCV*.
- Yang, C.; An, Z.; Cai, L.; and Xu, Y. 2021. Hierarchical Self-supervised Augmented Knowledge Distillation. *IJCAI*.
- Yim, J. 2017. A Gift from Knowledge Distillation: Fast Optimization, Network Minimization and Transfer Learning. *CVPR*.
- Zagoruyko, S.; and Komodakis, N. 2019. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*.
- Zbontar, J.; Jing, L.; Misra, I.; LeCun, Y.; and Deny, S. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. *ICML*.
- Zhang, L.; Song, J.; Gao, A.; Chen, J.; Bao, C.; and Ma, K. 2019. Be Your Own Teacher: Improve the Performance of Convolutional Neural Networks via Self Distillation.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful Image Colorization.
- Zhang, Y.; Lan, Z.; Dai, Y.; Zeng, F.; and Bai, Y. 2020. Prime-Aware Adaptive Distillation. *ECCV*.
- Zhang, Z.; and Sabuncu, M. R. 2020. Self-Distillation as Instance-Specific Label Smoothing.
- Zhao, B.; Song, R.; and Qiu, Y. 2022. Decoupled Knowledge Distillation. *CVPR*.
- Zhu, J.; Tang, S.; Chen, D.; and Yu, S. 2021a. Complementary Relation Contrastive Distillation. *CVPR*.
- Zhu, X.; Lyu, S.; Wang, X.; and Zhao, Q. 2021b. TPH-YOLOv5: Improved YOLOv5 Based on Transformer Prediction Head for Object Detection on Drone-captured Scenarios. *VisDrone ICCV workshop*.