Structural Knowledge Distillation for Object Detection

Anonymous Author(s)
Affiliation
Address
email

Abstract

Knowledge Distillation (KD) is a well-known training paradigm in deep neural networks where knowledge acquired by a large teacher model is transferred to a small student. KD has proven to be an effective technique to significantly improve the student's performance for various tasks including object detection. As such, KD techniques mostly rely on guidance at the intermediate feature level, which is typically implemented by minimizing an ℓ_p-norm distance between teacher and student activations during training. In this paper, we propose a replacement for the pixel-wise independent ℓ_p-norm based on the structural similarity (SSIM) [28]. By taking into account additional contrast and structural cues, more information within intermediate feature maps can be preserved. Extensive experiments on MSCOCO [17] demonstrate the effectiveness of our method across different training schemes and architectures. Our method adds only little computational overhead, is straightforward to implement and at the same time it significantly outperforms the standard ℓ_p-norms. Moreover, more complex state-of-the-art KD methods [14, 33] using attention-based sampling mechanisms are outperformed, including a +3.5 AP gain using a Faster R-CNN R-50 [22] compared to a vanilla model.

1 Introduction

Over the last decade, Deep Neural Networks (DNNs) have shown to be a very effective tool in solving fundamental computer vision tasks [15]. One major application of DNNs includes real-time perception systems found in e.g. autonomous vehicles, where object detection is often a task of major importance. Deployment of DNNs into real-time applications, however, introduces strict limitations on memory and latency. On the other hand, increased performance of state-of-the-art detectors typically comes with an increase in memory requirements and inference time [13]. Thus, the choice of network model and it's according detection performance is strictly limited. Several techniques have been proposed to tackle this problem, e.g. pruning [9], weight quantization [10], parameter prediction [6] and Knowledge Distillation (KD) [12]. In this work, we are particularly interested in the latter, as it provides an intuitive way of performance improvement without the need for architectural modifications to existing networks.

With KD, the knowledge acquired by a computationally expensive teacher model is transferred to a smaller student model during training. KD has proven to be very effective in tasks such as classification [12], segmentation [20], and in particular has seen considerable progress in detection very recently [5, 8, 14, 33, 34]. Due to the complexity of the output space of a typical detection model, it is necessary to apply KD at the intermediate feature level, as solely relying on output-based KD has proven ineffective [3, 5, 8, 14, 16, 25, 33, 34]. In feature-based KD, a training objective is introduced in addition to existing objectives, which minimizes the *error* between teacher and student activations

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

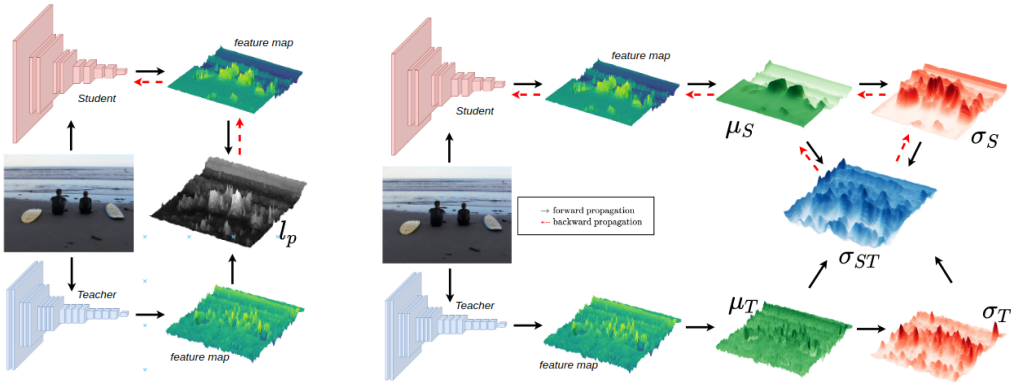

(a) Previous methods [25, 33, 34, 8, 5]. After sampling features an ℓ_p-norm is applied between feature maps.

(b) Our proposed method. We distill relational knowledge in the form of local mean μ, variance σ^2 and furthermore cross-correlation σ_{ST} between feature spaces.

Figure 1: Feature-based Knowledge Distillation (KD).

and is de-facto standard defined by the ℓ_p-norm distance [5, 8, 14, 25, 33, 34]. The ℓ_p-norm however ignores three important pieces of information present in the feature maps: (i) spatial relationships between features, (ii) the correlation between the teacher and student features and (iii) importance of individual features.

We notice recent work has (implicitly) focused on bypassing the latter point through the design of mechanisms that sample features by assuming that object regions are more "knowledge-dense" [5, 14, 25, 34]. However, as shown by [8], even distilling only background regions can improve performance, therefore it cannot be assumed that only the object regions contain useful knowledge. These sampling mechanisms furthermore introduce additional drawbacks which may limit their broader implementation into real-world applications, such as the need for labeled data [8, 14, 25].

In this work, we propose "Structural Knowledge Distillation", which attempts to improve the downsides associated with the ℓ_p-norm as a central driver for KD methods, rather than designing an ever more sophisticated sampling mechanism. Our key insight is that similar to images, the feature space can be locally decomposed into luminance, contrast and structure components, a strategy that has seen successful application in the image domain in the form of SSIM [28]. The new training objective becomes to minimize local difference in mean and variance, and maximize local zero-normalized cross-correlation between the teacher and student activations. We find that particularly the latter has great influence on distillation performance, as the underlying structure of relevant objects is propagated from the image throughout the CNN.

In order to demonstrate the effectiveness of our method we perform extensive experiments using various detection architectures and training schemes. Overall our contribution can be summarized as:

- We propose ℓ_{ssim} and variations as a replacement of the ℓ_p-norm for feature-based KD in object detection models, which enables the capture of knowledge manifested in contrast and structural cues.
- We demonstrate a consistent quantitative improvement in detection accuracy for various training settings and model architectures. With one line of code, ℓ_{ssim} outperforms purely feature based methods and performs on par or outperforms carefully tuned state-of-the-art object sampling mechanisms, without the associated complexity.
- We illustrate through an analysis of the feature space (i) that our method focuses on different areas than ℓ_p-norms and (ii) that simply sampling from foreground regions is suboptimal as background regions turn out to be knowledge-dense.

2 Related Work

Knowledge Distillation KD aims to transfer knowledge acquired by a cumbersome teacher model to a smaller student model. Bucilă et al. [1] (2006) demonstrate that the knowledge acquired by a large ensemble of models can be transferred to a single small model. Hinton et al. [12] (2015) provide a more general solution applied in a DNN in which they raise the temperature of the final softmax until the large model produces a suitably soft set of targets. Most KD research in the computer vision domain focuses on the classification task [24]. However, as our main interest lies in the real-time domain we focus on the more relevant object detection task.

Object Detection Object detection is one of the fundamental computer vision tasks, where speed and accuracy are often two key requirements. Object detectors can be classified into one-stage and two-stage methods, in this work we investigate our approach for both variations. The main meta-architecture within the one-stage domain is RetinaNet [18], with extensions including anchor-free modules and Reppoints [36, 32]. In the two-stage domain Faster R-CNN [22] is regarded as the most widely used meta-architecture, where iterations include Cascade R-CNN [2]. Furthermore, regardless of architecture, ResNet [11] backbones are often used to extract features, which are furthermore fused at multiple scales using e.g. a FPN [18].

Knowledge Distillation for Object Detection Several methods have been proposed that use KD for object detectors, where it has been found that typically guidance at the intermediate feature level rather than the output is critical due to the complex nature of the output space in detection models [3, 5, 8, 14, 16, 25, 33, 34]. As the detection task requires the identification of multiple objects at different locations, a major complexity introduced is the imbalance between foreground and background, which manifests itself in the intermediate features. Typically, the assumption is made that object regions are "knowledge-dense", and background regions less so. As a result, recent work has implicitly focused on designing mechanisms which sample object-relevant features to distill knowledge from [5, 14, 16, 25, 34].

Li et al. [16] (2017) mimic the features sampled from the region proposals in a two stage detector. Wang et al. [25] (2019) propose imitation masks which locate knowledge dense feature locations based on the annotated boxes. Dai et al. [5] (2021) propose a module which distills based on distance between classification scores. Similarly, Zhixing et al. [34] (2021) use output class probability to determine feature object probability. Recently Kang et al. [14] (2021) proposed a method in which they encode instance annotations in an attention mechanism [23] to locate "knowledge-dense" regions. Contrary to aforementioned methods, Zhang and Ma [33] (2021) propose a purely feature-based method in which they aim to both mimic the attention maps [35] as a sampling mechanism, and furthermore distill through non-local modules [26]. Regardless of the sampling technique, it has been demonstrated by [8] that it is not necessarily the case that background features are less important for distillation.

Objective Functions The objective in feature-based KD is to minimize the *error* between teacher and student feature spaces during training, typically in addition to existing objectives. The most widely used objective function in feature-based KD is the ℓ_p-norm with $p = 2$ [3, 5, 8, 16, 25, 33, 34], and less commonly $p = 1$ [3]. The ℓ_p-norm however, ignores the spatial relationships between features, the correlation between teacher and student and the importance of individual features, of which the latter has been the main focus of previous work. To take into account spatial dependencies, we need to furthermore compare features locally rather than pointwise. SSIM provides an elegant way to take into account spatial dependencies by making local comparisons of intensity and contrast, rather than just pointwise. It is further able to take into account the relationship between the teacher and the student by integrating zero-normalized cross-correlation. Contrary to alternative image signal quality metrics such as VIF [27], GMSD [27] and FSIM [31], SSIM is less complex to formulate mathematically and is differentiable, making it suitable as an objective function.

3 Method

3.1 Overview

We start off by defining the general form of feature-based distillation loss. For the purposes of this work, we divide a detector into three components: (i) the backbone, used for extracting features, (ii) the neck, for fusing features at different scales, and (iii) the head, for generating regression and classification scores. For feature-based KD, we select intermediate representations $\mathcal{T} \in \mathbb{R}^{C,H,W}$ and $\mathcal{S} \in \mathbb{R}^{C,H,W}$ from the teacher and student respectively at the output of the neck. The feature-based distillation loss between $\mathcal{T}$ and $\mathcal{S}$ can subsequently be formulated as:

$$\mathcal{L}_{feat} = \sum_{r=1}^{R} \frac{1}{N_r} \sum_{h=1}^{H} \sum_{w=1}^{W} \sum_{c=1}^{C} \mathcal{L}_{\varepsilon} \left(\nu \left(\phi \left(\mathcal{S}_{r,h,w,c} \right) \right), \nu \left(\mathcal{T}_{r,h,w,c} \right) \right) \tag{1}$$

where H, W, C, R are the height, width, number of channels and number of neck outputs respectively, $N_r = HWC$ the total number of elements for the r-th layer. Additionally we define $\nu(\cdot)$ as a *normalization* function which maps the values of $\mathcal{T}$ and $\mathcal{S}$ to $[0, 1]$, and $\phi(\cdot)$ as an optional *adaptation layer* [3] which matches the dimensionality of $\mathcal{T}$ and $\mathcal{S}$, in our case a 1×1 convolutional layer. We introduce the shortened notation $\mathcal{L}_{\varepsilon}$ which represents the choice of *difference* measurement function at a single feature position r, h, w, c on normalized features and including the adaptation layer, *i.e.* $\mathcal{L}_{\varepsilon} = \mathcal{L}_{\varepsilon} \left(\nu \left(\phi \left(\mathcal{S}_{r,h,w,c} \right) \right), \nu \left(\mathcal{T}_{r,h,w,c} \right) \right)$. Accordingly, we use $\mathcal{S}$ and $\mathcal{T}$ to denote normalized and adapted student and normalized teacher activations respectively, e.g. $\mathcal{S} = \nu \left(\phi \left(\mathcal{S}_{r,h,w,c} \right) \right)$.

3.2 Measuring Difference

As ascertained, the de-facto standard choice for $\mathcal{L}_{\varepsilon}$ is the ℓ_p norm. $p = 2$ penalizes large errors, but is more tolerable to smaller errors. On the other hand, $p = 1$ does not over-penalize large errors, but smaller errors are penalized more harshly. The ℓ_p norm in its general form is given by:

$$\ell_p : \mathcal{L}_{\varepsilon} = \left(|\mathcal{S} - \mathcal{T}|^p \right)^{1/p} \tag{2}$$

Clearly such a function is not able to capture spatial relationships between features. In order to capture second-order information we need to involve at least two feature positions, we therefore change the problem statement from a *point-wise* comparison to a local *patch-wise* comparison. For each such patch, we extract three fundamental properties: the mean μ, the variance σ^2, and the cross-correlation $\sigma_{\mathcal{ST}}$ which captures the relationship between $\mathcal{S}$ and $\mathcal{T}$. We follow [28] and compute these quantities using a Gaussian-weighted patch F_{σ_F} of size 11×11 and $\sigma_F = 1.5$. The proposed SSIM framework [28] compares each of the properties, and is therefore composed of three components: luminance l, contrast c and structure s, which are defined as follows:

$$l = \frac{2\mu_{\mathcal{S}}\mu_{\mathcal{T}} + C_1}{\mu_{\mathcal{S}}^2 + \mu_{\mathcal{T}}^2 + C_1} \tag{3a} \qquad c = \frac{2\sigma_{\mathcal{S}}\sigma_{\mathcal{T}} + C_2}{\sigma_{\mathcal{S}}^2 + \sigma_{\mathcal{T}}^2 + C_2} \tag{3b} \qquad s = \frac{\sigma_{\mathcal{ST}} + C_3}{\sigma_{\mathcal{S}}\sigma_{\mathcal{T}} + C_3} \tag{3c}$$

where $\mu_{\mathcal{S}}$, $\mu_{\mathcal{T}}$ refer to the mean, $\sigma_{\mathcal{S}}$, $\sigma_{\mathcal{T}}$ refer to the variance and $\sigma_{\mathcal{ST}}$ refers to the covariance within the patch. Furthermore, to prevent instability $C_1 = (K_1 L)^2$, $C_2 = (K_2 L)^2$, $C_3 = C_2/2$, where L is the dynamic range of the feature map and $K_1 = 0.01$, $K_2 = 0.03$. An important property of SSIM is that it assigns more importance to *relative* changes in l and c due to the quadratic terms in the denominator. Furthermore, s is a direct measurement of the zero-normalized correlation coefficient between $\mathcal{S}$ and $\mathcal{T}$, and hence is formulated as the ratio between their covariance and product of standard deviations. As the range of SSIM is $[-1, 1]$, combining these three components results in the following objective:

$$\ell_{\text{SSIM}} : \mathcal{L}_{\varepsilon} = (1 - \text{SSIM})/2 = (1 - \left(l^{\alpha} \cdot c^{\beta} \cdot s^{\gamma} \right))/2 \,, \tag{4}$$

where the prevalence of each function can be tuned, with $\alpha = \beta = \gamma = 1.0$ as a default. As our method is purely feature-based and therefore independent of the type of head or bounding box labels, we simply add $\mathcal{L}_{feat}$ to the existing detection objective function $\mathcal{L}_{det}$ (typically $\mathcal{L}_{cls}$ and $\mathcal{L}_{reg}$) using weighting factor λ, which results in the following overall training objective:

$$\mathcal{L} = \lambda \mathcal{L}_{feat} + \mathcal{L}_{det} \tag{5}$$

Note that the implementation of ℓ_{SSIM} only requires the replacement of one line of code compared to ℓ_p, an example is provided in the supplemental material.

4 Experiments

4.1 Experiment Settings

Following literature [8, 14, 25, 33, 34], we assess the performance of ℓ_{SSIM} on the MSCOCO [17] validation dataset. We report mean Average Precision (AP) as the main evaluation metric, and additionally report AP at specified IoU thresholds AP_{50}, AP_{75} and object sizes AP_S, AP_M, AP_L. Our central points of comparison are the two most widely used one- and two-stage meta-architectures, RetinaNet (RN) [19] and Faster-RCNN (FRCNN) [22]. We respectively use ResNet/ResNeXt-101 backbones [11, 30] for the teachers and R-50 [11] backbones for all students, with $\lambda = 4, 2$ respectively. We conduct our experiments in Pytorch [21] using the MMDetection2 [4] framework on a Nvidia RTX8000 GPU with 48GB of memory. Each model is trained using SGD optimization with momentum 0.9, weight decay 1e-4 and batch size 8. The learning rate is set at 0.01 (RN) / 0.02 (FRCNN) and decreased tenfold at step 8 and 11, for a total of 12 epochs.

4.2 Comparison with ℓ_p-norms

In this first set of experiments we compare the performance of ℓ_p-norms to ℓ_{SSIM}. Table 1 shows the results of the main experiments comparing the best performance of each $\mathcal{L}_\varepsilon$. It can be observed that: (i) ℓ_{SSIM} outperforms ℓ_p-norms by a significant margin, boosting performance with up to +3.7AP. (ii) Adopting any form of feature-based distillation results in an improvement over the vanilla network, except in Faster R-CNN [22]. (iii) Even though previous work uses ℓ_2, ℓ_1 outperforms ℓ_2 with AP improvements of 2.3 vs. 0.4 and 1.2 vs. 0.0 respectively.

Table 1: Comparison of objective functions on MSCOCO [17].

Backbone	$\mathcal{L}_\varepsilon$	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
RetinaNet [19]							
Teacher R101		41.0	60.3	44.0	24.1	45.3	53.8
Vanilla R50	-	*36.4*	*55.6*	*38.7*	*21.1*	*40.3*	*46.6*
R50	ℓ_2	36.8 (+0.4)	55.7	39.1	20.6	40.5	47.3
R50	ℓ_1	38.7 (+2.3)	57.6	41.6	22.7	42.7	50.5
R50	$\boldsymbol{\ell}_{\text{SSIM}}$	**40.1 (+3.7)**	**59.2**	**43.1**	**23.1**	**44.6**	**53.2**
Faster R-CNN R50 [22]							
Teacher X-101		45.6	64.1	49.7	26.2	49.6	60.0
Vanilla R50		*37.4*	*58.1*	*40.4*	*21.2*	*41.0*	*48.1*
R50	ℓ_2	37.4 (+0.0)	57.6	40.9	21.2	41.3	48.1
R50	ℓ_1	38.6 (+1.2)	58.8	42.1	21.8	42.1	49.9
R-50	$\boldsymbol{\ell}_{\text{SSIM}}$	**40.9 (+3.5)**	**61.0**	**44.9**	**23.7**	**44.5**	**53.5**

To investigate what this improvement in performance can be attributed to, we analyze the distribution of the training stimulus in the feature space. Figure 2 illustrates the comparison between ℓ_{ssim} and ℓ_p ($p = 2$) for the magnitude of the loss, averaged over all channels and 12 training epochs in neck $r = 1$. This tells us something about which regions in the feature space are focused on more. It can be observed that with ℓ_2, high loss is assigned to object regions in particular, and furthermore regions with high brightness, such as the sky in fig. 2a or the window in fig. 2b. ℓ_{ssim} however assigns the loss differently, where not only object regions are focused, but additionally more diverse background regions are targeted, while little importance is given to low-contrast background regions.

One of the issues highlighted by [8] is that losses are higher in object regions than in background regions. As can be seen in fig. 2, the loss applied by ℓ_{ssim} is much more distributed over the feature

(a) (b) (c) (d)

Figure 2: Distribution of the magnitude of the loss in the feature space $r = 1$ for various images, averaged over all channels and 12 epochs of training on MSCOCO [17]. From left to right: Image, student trained with ℓ_2, student trained with ℓ_{ssim}. A darker color indicates a higher loss, object regions have been highlighted with bounding boxes, and feature maps have been normalized.

space than ℓ_2, which as a direct result causes a more distributed application of the gradient in the feature space. As a result, is can be observed that the feature map of a ℓ_{ssim} distilled model is much more similar to the teacher than an ℓ_p distilled model, as shown in fig. 3, which directly translates to the increase in performance.

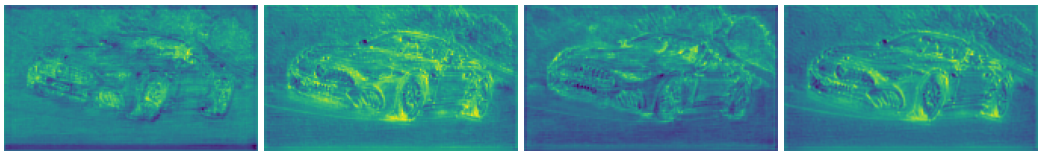

(a) Vanilla (b) Teacher (c) ℓ_2 (d) ℓ_{ssim}

Figure 3: Qualitative comparison of channel sampled randomly from RetinaNet [19] intermediate neck features $r = 1$. Lighter colors indicate higher activations.

4.3 Influence of Luminance, Contrast and Structure

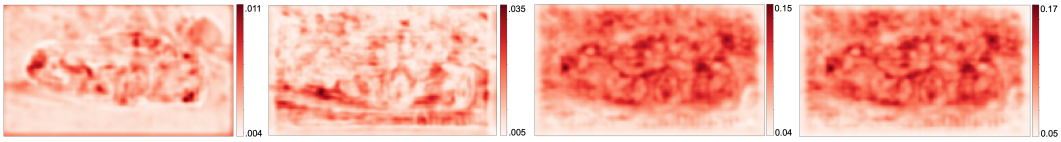

(a) Luminance (α) (b) Contrast (β) (c) Structure (γ) (d) ssim (α, β, γ)

Figure 4: Distribution of the channel averaged magnitude of the loss for each individual component luminance, contrast and stucture, averaged over 12 epochs of training on MSCOCO [17].

Next we compare the influence of the luminance, contrast and structure components by tuning α, β and γ respectively (refer to eq.(4)), for this experiment we use the RetinaNet [19] teacher-student pair. The results are shown in table 2. The main observation that can be made here is that the influence of structure (γ) is more substantial than the other components, and even on its own provides an increase of +3.2 AP. Furthermore, the comparisons of luminance α and contrast β alone result in performance comparable or better than ℓ_p-norms (compare to table 1). In fig. 4 we again compare the magnitude of the loss over the entire training period, where it can be observed that structure has the highest overall impact. The luminance provides a similar training stimulus to the ℓ_p-norm (fig. 2c), but is more "smooth" due to the Gaussian kernel. Contrarily, contrast mostly targets background areas. Finally, it can be noticed that structure has the most influence over the total loss, both in magnitude and spatial distribution.

Table 2: Comparison of objective functions for RetinaNet [19] on MSCOCO [17]. α tunes luminance, β tunes contrast and γ tunes structure.

Backbone	α	β	γ	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
Teacher ResNet-101				41.0	60.3	44.0	24.1	45.3	53.8
Vanilla ResNet-50				*36.4*	*55.6*	*38.7*	*21.1*	*40.3*	*46.6*
ResNet-50	1	0	0	38.7 (+2.3)	57.9	41.6	21.8	42.8	50.8
ResNet-50	0	1	0	38.9 (+2.5)	57.7	41.6	21.7	42.6	51.3
ResNet-50	0	0	1	39.6 (+3.2)	58.6	42.7	22.5	44.0	52.5
ResNet-50	0	1	1	40.0 (+3.6)	59.0	42.8	22.4	44.4	**53.3**
ResNet-50	1	1	1	**40.1 (+3.7)**	**59.2**	**43.1**	**23.1**	**44.6**	53.2

In fig. 5 we illustrate the differences between student and teacher activations for the differently trained models. It can be observed that the structure objective results in a feature space that is convered to a very similar local optimum as the teacher, with few noisy or large differences. Luminance contains more noisy differences, and especially in the last layer demonstrated high differences. Although the contrast objective performs relatively well on its own, it seems to provide different activations as the teacher, particularly in the object area, as it converged to a different local minimum.

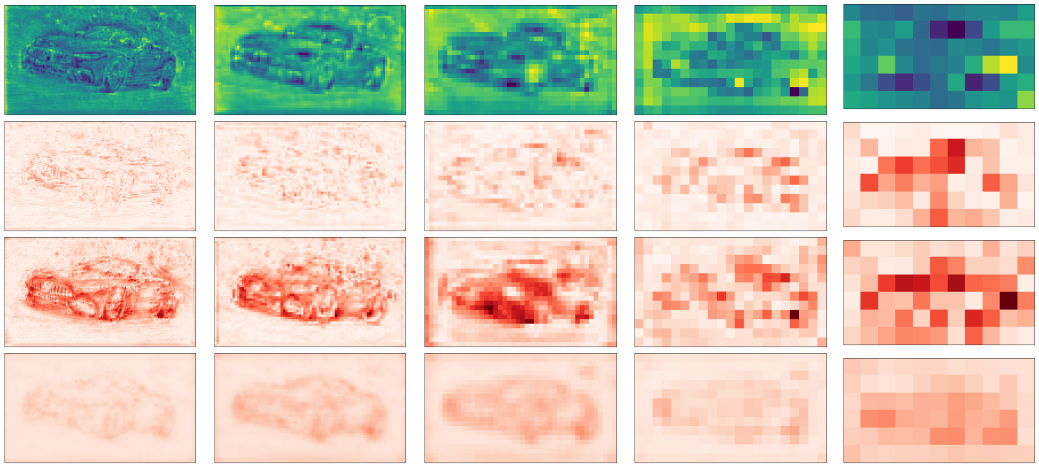

(a) $r = 1$ (b) $r = 2$ (c) $r = 3$ (d) $r = 4$ (e) $r = 5$

Figure 5: Top row: Channel averaged activations in the RetinaNet R101 [19] Teacher. Subsequent rows illustrate channel averaged differences in activations between teacher and students, distilled with: 2nd row: luminance (α). 3rd row: contrast (β). 4th row: structure (γ). Differences have been normalized, where darker color indicates a higher value.

4.4 Comparison to State-of-the-Art

Next, we compare to recent work, for which the following methods serve as baselines: (i) Zhang and Ma [33] (2021), a purely feature-based approach leveraging attention masks [35] and non-local modules [26], and (ii) Kang et al. [14] (2021), who encode labeled instance annotations in an attention mechanism [23] and report state-of-the-art for distillation methods for RetinaNet [19] and Faster R-CNN [22] on MSCOCO [17] at the time of writing. To further demonstrate the simplicity and versatility of our method, we use the original code and teachers as the authors to compare to our proposed method in the same experimental setup. [33] use the same MMDetection2 [4] framework, while [14] use Detectron2 [29]. For the comparison with [14] we furthermore adopt inheritance, a practice proposed by the authors in which the FPN [18] and head of the student are initialized with teacher parameters. This leads to faster training convergence, but may not be applicable when

architectures differ between teacher and student. As the teachers and exact configurations slightly vary, we split up the comparison into two parts, as shown in table 3.

Table 3: Comparison to state-of-the-art methods on MSCOCO [17]. † denotes inheritance.

Method	RetinaNet [19]				Faster R-CNN [22]			
	AP	AP_S	AP_M	AP_L	AP	AP_S	AP_M	AP_L
Teacher	41.0	24.1	45.3	53.8	45.6	26.2	49.6	60.0
Vanilla	*36.4*	*21.1*	*40.3*	*46.6*	*37.4*	*21.2*	*41.0*	*48.1*
Zhang and Ma [33]	38.5 (+2.1)	21.7	42.6	51.5	38.9 (+1.5)	21.9	42.1	51.5
Ours	**40.1 (+3.7)**	**23.1**	**44.6**	**53.2**	**40.9 (+3.5)**	**23.7**	**44.5**	**53.5**
Teacher	40.4	24.0	44.3	52.2	42.0	25.2	45.6	54.6
Vanilla	*37.4*	*23.1*	*41.6*	*48.3*	*37.9*	*22.4*	*41.1*	*49.1*
Kang et al. [14] †	**40.7 (+3.3)**	**24.6**	44.9	52.4	40.9 (+3.0)	**24.5**	44.2	53.3
Ours †	**40.7 (+3.3)**	24.0	**45.0**	**53.1**	**41.0 (+3.1)**	23.8	**44.5**	**53.7**

It can be observed that: (i) the adoption of SSIM as the distillation function results in an improvement of +3.7 AP, and outperforms [33] for all box sizes and IoU thresholds. (ii) Both SSIM and [14] result in an improvement of +3.3 AP over the vanilla network. In particular, our method scores high for AP_L, while [14] mainly show better performance in the small object AP_S category. Additionally it can be observed that the student is able to outperform the teacher with RetinaNet [19].

4.5 Ablation Studies

Generalizability to Detection Architectures and Schedules We perform additional studies on several different detection architectures to demonstrate the generalizability of our method. We evaluate our distillation method on the smaller ResNet-18 backbone and two alternative one-stage architectures, Fsaf-RetinaNet [36], which extends RetinaNet [19] with an anchor-free module, and Reppoints [32], which replaces the regular bounding box representation of objects by a set of sample points. The results of our experiments are shown in Table 4. It can be observed that: (i) For each detection architecture our method significantly improves performance, with +3.5AP for the ResNet18 backbone, +2.3 AP for Fsaf-RetinaNet [36], +3.3 AP for Reppoints [32]. (ii) In general, our method is modular and can significantly improve performance regardless of the detection framework used.

Table 4: Investigation of several popular detection architectures on MSCOCO [17].

Model	AP	AP_{50}	AP_{75}	AP_S	AP_M	AP_L
RetinaNet-R101 (Teacher)	41.0	60.3	44.0	24.1	45.3	53.8
RetinaNet-R18 (Vanilla)	32.6	50.6	34.6	17.8	35.2	43.5
RetinaNet-R18 (Ours)	**36.1 (+3.5)**	**54.3**	**38.6**	**18.9**	**39.7**	**49.2**
RetinaNet-R101 (Teacher)	41.0	60.3	44.0	24.1	45.3	53.8
RetinaNet-R50 (Vanilla, 2x)	37.4	56.7	39.6	20.0	40.7	49.7
RetinaNet-R50 (Ours, 2x)	**40.6 (+3.2)**	**59.7**	**43.7**	**23.6**	**44.8**	**53.9**
Fsaf-RetinaNeXt-X101 [36] (Teacher)	42.4	62.5	45.5	24.6	46.1	55.5
Fsaf-RetinaNet-R50 [36] (Vanilla)	37.4	56.8	39.8	20.4	41.1	48.8
Fsaf-RetinaNet-R50 [36] (Ours)	**39.7 (+2.3)**	**59.3**	**42.4**	**22.0**	**43.3**	**52.0**
Reppoints X-101[32] (Teacher)	44.2	65.5	47.8	26.2	48.4	58.5
Reppoints-R50 [32] (Vanilla)	37.0	56.7	39.7	20.4	41.0	49.0
Reppoints-R50 [32] (Ours)	**40.3 (+3.3)**	**60.3**	**43.5**	**22.6**	**44.4**	**53.9**

Effects of an Adaptation Layer Adaptation layers can be implemented when channel or spatial dimensions between teacher and student do not match, and have shown to generally improve performance previous methods [3, 25, 33]. We adapt the commonly used 1×1 convolution to investigate

the influence on performance with our method. We adopt the RetinaNet R101-50 teacher-student pair and the CRCNN X101 - FRCNN R50 teacher-student pair. The results are shown in table 5. We notice that in table 5a there is no additional benefit of adopting an adaptation layer, while in table 5b the difference is significant, and implementing the adaptation layer is critical.

Table 5: Investigation of the effect of adaptation layers

(a) RetinaNet R101 - R50 [19]

Adap. layer	AP	AP_S	AP_M	AP_L
none	**40.1**	**23.1**	**44.6**	53.2
1×1	**40.1**	**23.1**	44.4	**53.4**

(b) Cascade RCNN X101 [2] - FRCNN R50 [22]

Adap. layer	AP	AP_S	AP_M	AP_L
none	39.8	22.6	43.4	52.1
1×1	**40.9**	**23.7**	**44.5**	**53.5**

Varying Patch Size F and loss prevalence λ We additionally investigate the two remaining main hyperparameters that we introduce in this work, for which we use the RetinaNet R101-50 [19] teacher-student pair. The influence of the prevalence of $\mathcal{L}_{feat}$ tuned by λ is shown in fig. 6, where it can be noticed that the choice of λ can cause a difference of up to +0.5 AP, with $\lambda = 2, 4$ providing the best performance. Additionally the influence of the local patch size F over which we calculate each component of ℓ_{ssim} is investigated, the results are shown in fig. 6. It can be noticed that the choice of kernel size does not have significant influence over performance.

Error types Finally we are interested in the type of improvements made by our distillation method, as illustrated in fig. 7. We can particularly see improvements in the AP_L category, which can also be noticed in table 3 where we perform better in this category than previous work. One explanation is illustrated in fig. 5d, 5e where we can see that the structural part ensures convergence towards the teacher, particularly in the deeper layers which are responsible for detecting large objects. We furthermore observe in particular improvements in localization, as seen in the AP_{75}, AP_{50} and AP_{10} (Loc) categories in fig. 7.

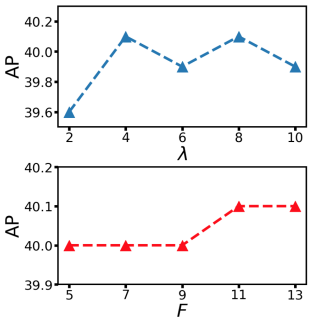

Figure 6: Top: varying λ. Bottom: varying F.

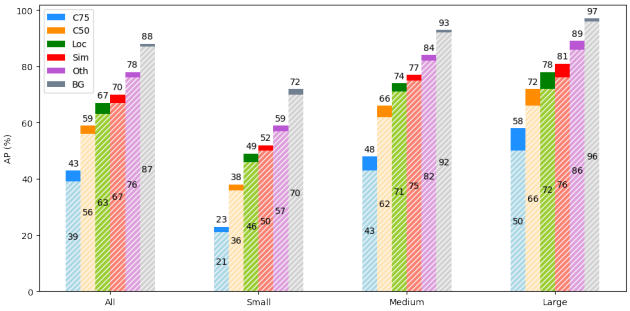

Figure 7: AP score for varying box sizes. Hatched areas represent the baseline, solid areas represent the performance increase obtained through distillation.

5 Conclusion

This paper proposed ℓ_{ssim} as a replacement for the conventional ℓ_p-norm as a building block for feature-based KD in object detection, which allows us to take into account additional contrast and structural cues which improve the optimization of the student. Through extensive experiments we show that even on its own ℓ_{ssim} is able to reach performance on par or even outperform state-of-the-art, without the need for carefully designed sampling mechanisms which add complexity and need to be tuned during training. Our method is simple and can be implemented by replacing one line of code. Future work consists of using ℓ_{ssim} as a building block for future KD methods, which includes applying it to other computer vision tasks.

References

[1] C. Bucilă, R. Caruana, and A. Niculescu-Mizil. Model compression. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006:535–541, 2006. doi: 10.1145/1150402.1150464.

[2] Z. Cai and N. Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

[3] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker. Learning efficient object detection models with knowledge distillation, 2017. ISSN 10495258.

[4] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin. {MMDetection}: Open MMLab Detection Toolbox and Benchmark. *arXiv preprint arXiv:1906.07155*, 2019.

[5] X. Dai, Z. Jiang, Z. Wu, Y. Bao, Z. Wang, S. Liu, and E. Zhou. General Instance Distillation for Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7842–7851, 2021.

[6] M. Denil, B. Shakibi, L. Dinh, M. Ranzato, and N. De Freitas. Predicting parameters in deep learning. *Advances in Neural Information Processing Systems*, pages 1–9, 2013. ISSN 10495258.

[7] D. P. E. R. E. Riba D. Mishkin and G. Bradski. Kornia: an Open Source Differentiable Computer Vision Library for PyTorch. In *Winter Conference on Applications of Computer Vision*, 2020. URL `https://arxiv.org/pdf/1910.02190.pdf`.

[8] J. Guo, K. Han, Y. Wang, H. Wu, X. Chen, C. Xu, and C. Xu. Distilling object detectors via decoupled features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2154–2164, 2021.

[9] S. Han, J. Pool, J. Tran, and W. J. Dally. Learning both weights and connections for efficient neural networks. *Advances in Neural Information Processing Systems*, 2015-Janua:1135–1143, 2015. ISSN 10495258.

[10] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, pages 1–14, 2016.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[12] G. Hinton, O. Vinyals, and J. Dean. Distilling the Knowledge in a Neural Network. pages 1–9, 2015. URL `http://arxiv.org/abs/1503.02531`.

[13] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and others. Speed/accuracy trade-offs for modern convolutional object detectors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7310–7311, 2017.

[14] Z. Kang, P. Zhang, X. Zhang, J. Sun, and N. Zheng. Instance-Conditional Knowledge Distillation for Object Detection. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.

[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

[16] Q. Li, S. Jin, and J. Yan. Mimicking very efficient network for object detection. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-Janua: 7341–7349, 2017. doi: 10.1109/CVPR.2017.776.

[17] T. Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8693 LNCS(PART 5):740–755, 2014. ISSN 16113349. doi: 10.1007/978-3-319-10602-1{_}48.

[18] T. Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-Janua:936–944, 2017. doi: 10.1109/CVPR.2017.106.

[19] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.

[20] Y. Liu, K. Chen, C. Liu, Z. Qin, Z. Luo, and J. Wang. Structured knowledge distillation for semantic segmentation. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2019-June:2599–2608, 2019. ISSN 10636919. doi: 10.1109/CVPR.2019.00271.

[21] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d\textquotesingle Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

[22] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. 2015. URL `http://arxiv.org/abs/1506.01497`.

[23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.

[24] L. Wang and K. J. Yoon. Knowledge Distillation and Student-Teacher Learning for Visual Intelligence: A Review and New Outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(8):1–40, 2021. ISSN 19393539. doi: 10.1109/TPAMI.2021.3055564.

[25] T. Wang, L. Yuan, X. Zhang, and J. Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019.

[26] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.

[27] Z. Wang and Q. Li. Information content weighting for perceptual image quality assessment. *IEEE Transactions on image processing*, 20(5):1185–1198, 2010.

[28] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[29] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick. Detectron2. \url{https://github.com/facebookresearch/detectron2}, 2019.

[30] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.

[31] W. Xue, L. Zhang, X. Mou, and A. C. Bovik. Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE transactions on image processing*, 23(2): 684–695, 2013.

[32] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9657–9666, 2019.

[33] L. Zhang and K. Ma. Improve Object Detection with Feature-based Knowledge Distillation: Towards Accurate and Efficient Detectors. In *International Conference on Learning Representations*, 2021.

[34] D. Zhixing, R. Zhang, M. Chang, S. Liu, T. Chen, Y. Chen, and others. Distilling Object Detectors with Feature Richness. *Advances in Neural Information Processing Systems*, 35, 2021.

[35] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.

[36] C. Zhu, Y. He, and M. Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 840–849, 2019.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Because of resource limitations/training hardware, and it is furthermore not commonplace in this domain.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] Yes, we only use publicly available data and code: MSCOCO [17], PyTorch [21], MMDetection2 [4], Detectron2 Wu et al. [29], Kornia [7].
   (b) Did you mention the license of the assets? [No] Citing is sufficient according to the authors.
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]