

G-DetKD: Towards General Distillation Framework for Object Detectors via Contrastive and Semantic-guided Feature Imitation

Lewei Yao^{1*} Renjie Pi^{1*} Hang Xu^{2†} Wei Zhang² Zhenguo Li² Tong Zhang¹
¹Hong Kong University of Science and Technology ²Huawei Noah's Ark Lab

In this paper, we investigate the knowledge distillation (KD) strategy for object detection and propose an effective framework applicable to both homogeneous and heterogeneous student-teacher pairs. The conventional feature imitation paradigm introduces imitation masks to focus on informative foreground areas while excluding the background noises. However, we find that those methods fail to fully utilize the semantic information in all feature pyramid levels, which leads to inefficiency for knowledge distillation between FPN-based detectors. To this end, we propose a novel semantic-guided feature imitation technique, which automatically performs soft matching between feature pairs across all pyramid levels to provide the optimal guidance to the student. To push the envelop even further, we introduce contrastive distillation to effectively capture the information encoded in the relationship between different feature regions. Finally, we propose a generalized detection KD pipeline, which is capable of distilling both homogeneous and heterogeneous detector pairs. Our method consistently outperforms the existing detection KD techniques, and works when (1) components in the framework are used separately and in conjunction; (2) for both homogeneous and heterogeneous student-teacher pairs and (3) on multiple detection benchmarks. With a powerful X101-FasterRCNN-Instaboost detector as the teacher, R50-FasterRCNN reaches 44.0% AP, R50-RetinaNet reaches 43.3% AP and R50-FCOS reaches 43.1% AP on COCO dataset.

1. Introduction

Knowledge distillation (KD) is a training strategy aiming at transferring the learnt knowledge from a powerful, cumbersome model (teacher) to a more compact model (student). The seminal work [11] introduced the idea of KD and this technique has been proven to be effective on classification tasks by many subsequent works [36, 8, 26]. However, integrating the idea of KD into detection is nontrivial. The conventional paradigm of classification KD can

*Equal contribution

†Corresponding author: xbjxh@live.com



Figure 1. Our G-DetKD consistently improves the performance of detectors belonging to various categories by around 4% in AP.

not be directly applied to detection task since simply minimizing the KL divergence between the classification outputs fails to extract the spatial information from the teacher and only brings limited performance gain [29] to the student. In this work, we aim at developing a general KD framework which can efficiently extract the spatial information and is applicable to both homogeneous and heterogeneous student-teacher pairs.

It is acknowledged that feature map imitation with foreground attention mechanisms helps students learn better [2, 14, 29, 25]. Previous works propose different imitation masks to focus on the informative foreground regions while excluding the background noises. However, mask-based methods were first developed for outdated detectors, i.e., vanilla Faster-RCNN without FPN [23], which fail to extend to modern detectors equipped with FPN. Specifically, those methods perform direct one-to-one matching between pyramid levels of the student-teacher pair, which leads to two issues: (1) indiscriminately applying the same mask on all levels can introduce noise from unresponsive feature levels; (2) mask-based methods are not extendible to heterogeneous detector pairs since their feature levels may not be strictly aligned, e.g., FasterRCNN constructs the feature pyramid from P_2 to P_6 , while RetinaNet uses P_3 to P_7 .

To address the above issues, we propose a simple yet effective Semantic-Guided Feature Imitation (SGFI) approach based on object proposals. By analyzing the re-

response patterns of feature pyramids as shown in Figure 2, we find that the best matching features for imitation may be from different pyramid levels for the student-teacher pair. In addition, the object regions on nearby feature levels are activated in a similar pattern, which implies that student features from other levels carrying similar semantics can also benefit from imitation. Thus, to fully exploit the potential of teacher’s feature pyramid, our proposed method automatically performs soft matching between the features based on their semantic relationship to provide the student with the optimal guidance. The proposed SGFI outperforms other feature imitation approaches by a large margin.

Feature imitation techniques mimic the features corresponding to the same region, whereas the relation between representations of different regions also encodes informative knowledge to help the student’s learning. Thus, we further propose Contrastive Knowledge Distillation (CKD) for detection inspired by the idea of contrastive learning. Specifically, we make use of the proposal region’s representations produced by the student-teacher pair to construct positive and negative contrastive pairs, then an InfoNCE loss is adopted to minimize the distance between the positive pairs while pushing away the negative ones. We show that the standalone contrastive approach demonstrates outstanding effectiveness and can further boost the detector’s performance when applied with our SGFI.

In some scenarios, only detectors with certain architectures can be deployed due to hardware constraints, while the most powerful teachers belong to different categories. In this case, knowledge distillation between heterogeneous detector pairs is promising. However, previous works only consider KD between homogeneous detectors pairs due to their lack of extensibility. Therefore, we are motivated to extend the two approaches mentioned above into a general detection KD framework called **G-DetKD** which is applicable to both homogeneous and heterogeneous student-teacher detector pairs.

Extensive experiments demonstrate the effectiveness of our proposed G-DetKD framework. On COCO dataset [17], without bells and whistles, G-DetKD easily achieves state-of-the-art performance among detection KD methods. **As no modification is applied to the student detectors, the performance gain is totally free.** Using the same two-stage detector as the teacher, the performance gains for homogeneous students surpass other SOTA detection KD methods: R50-FasterRCNN reaches 44.0% AP; while the effect on heterogeneous students is also significant: R50-Retina reaches 43.3 % AP and R50-FCOS reaches 43.1% AP. Our method generalizes surprisingly well for large detectors like CascadeRCNN with ResNeXt101-DCN as the backbone: boosting its AP from 46.0% to 50.5%. In addition, the generalization ability of our method is validated on multiple mainstream detection benchmarks, e.g., Pascal

VOC [6] and BDD [34].

In summary, the contributions of this paper are threefold:

- We propose a novel semantic-guided feature imitation approach (SGFI) with a semantic-aware soft-matching mechanism.
- We propose contrastive knowledge distillation (CKD) to capture the information encoded in the relationship between teacher’s different feature regions.
- We make the first attempt to construct a general KD framework (G-DetKD) capable of distilling knowledge for both homogeneous and heterogeneous detectors pairs. Comprehensive experiments are conducted to show the significant performance boosts brought by our approach.

2. Related Works

Object Detection. Object detection is one of the fundamental problems in computer vision. State-of-the-art detection networks can be categorized into one-stage, two-stage and anchor-free detectors. One-stage detectors such as [21, 18, 22] perform object classification and bounding box regression directly on feature maps. On the other hand, two-stage detectors such as [23, 15] adopt a “coarse-to-fine” approach which uses a region proposal network (RPN) to separate foreground boxes from background and a RCNN head to refine the regression results and perform the final classification. [28, 5, 31] propose to directly predict location of objects rather than based on anchor priors, which opens a new era for object detection. Recent works also perform NAS on detection tasks, which searches for novel detectors automatically without human intervention [12, 30, 9, 32, 33].

Knowledge Distillation. KD was first proposed in [11]. Its effectiveness has been explored by many subsequent works [36, 8, 26, 24]. For object detection, [2] first proposes imitating multiple components in detection pipeline, e.g., backbone features, RPN and RCNN. Recent works demonstrate that foreground instances are more important in feature imitation. Various methods are proposed to help student model focus on foreground information, including multiple mask generating approaches [27, 25] and RoI extraction [14]. However, these methods are designed for detectors without FPN, which fail to extend to FPN-based detectors with multiple feature levels. This motivates us to design a new KD framework that is able to fully exploit teacher’s feature pyramid.

Contrastive Learning. Contrastive learning is a popular approach for self-supervised tasks [19, 10, 4]. The goal is to bring closer the representations of similar inputs and push away those of dissimilar ones, which naturally takes into account the relationship between contrastive pairs. Inspired by the recent work [27], which proposes a distilla-

tion method using a contrastive approach for classification, we propose to integrate the information encoded in the relationship between different object regions by introducing contrastive knowledge distillation for object detection.

3. Preliminary

We start by briefly reviewing previous works on KD for object detection. [2] conducted knowledge distillation on detector’s classification and localization predictions, which are formulated as: $L_{cls} = -\frac{1}{N} \sum^N \mathbf{P}_t \log \mathbf{P}_s$, $L_{reg} = -\frac{1}{N} \sum^N |reg_t^i - reg_s^i|$, where P_s , P_t and reg_s , reg_t are the class scores and localization outputs of the student-teacher pair, respectively. However, simply distilling from the prediction outputs neglects the teacher’s spatial information.

More recent works mainly focus on distilling knowledge from the feature maps, which encode spatial information that is crucial for detection. These methods propose different imitation masks I to form an attention mechanism for foreground features and filters away excessive background noises. The objective can be formulated as:

$$L_{feat} = \frac{1}{2N_p} \sum_{l=1}^L \sum_{i=1}^W \sum_{j=1}^H \sum_{c=1}^C I_{ij}^l \left(f_{adapt}^l(S^l)_{ijc} - T_{ijc}^l \right)^2$$

where L , W , H , C are feature levels, width, height and number of channels, respectively; S and T are student’s and teacher’s feature maps; I is the imitation mask; $f_{adapt}(\cdot)$ is an adaptation function mapping S and T to the same dimension, which is usually implemented as a 3x3 conv layer; N_p is the total number of positive elements in the mask.

Various methods can differ in the definition of I . For example, FGFI [29] generates an imitation mask according to the near object anchor locations, while TADF [25] proposed to generate a soft mask using Gaussian function.

Those mask-based methods are originally proposed for detectors without FPN. To extend them for the FPN-based modern detectors, the masks are generated by the same rule for all feature levels, then features on the corresponding levels are matched for imitation. We argue this direct adaptation to modern detectors with feature pyramids is suboptimal due to the following reasons: (1) generally, in the design paradigm of FPN, each feature level is responsible for detecting objects of different scales. Thus, indiscriminately applying the same mask on all levels can introduce noise from unresponsive feature levels; (2) mask-based methods are not extendible to heterogeneous detector pairs since their feature pyramid levels may not be strictly aligned, e.g., Faster-RCNN-FPN constructs the feature pyramid from P_2 to P_6 , while RetinaNet-FPN uses P_3 to P_7 . These weaknesses promote us to design a feature imitation mechanism that can automatically match the features of the student-




Figure 2. Activation patterns of feature pyramids from different detectors. We observe that: (1) the same feature level from different detectors present various patterns; (2) nearby feature maps contain similar semantics.

teacher pair for imitation and eliminate the excessive noise, while also being extendible to heterogeneous detector pairs.

4. Methods

4.1. Semantic-Guided Feature Imitation (SGFI)

The aforementioned problem of mask-based methods can be partially solved via a straightforward approach which uses positive RoI features for imitation: instead of uniformly imitating foreground regions on all feature levels, RoI extraction operation selects foreground features based on object proposals in a fine-grained manner, which automatically matches the features from corresponding levels according to their scales. However, RoI extractor’s proposal assignment heuristics is purely based on the proposal scales, which is agnostic to the semantic relationship between features of the student-teacher pair.

Intuitively, a promising feature distillation approach needs to consider the semantics of features when constructing them into pairs for imitation. We reflect on the feature matching mechanism by looking into the characteristics of feature pyramids, which are visualized in Figure 2. We observe that for different detectors, given an object in the image, the corresponding region on the same feature level present different patterns, while the difference is more significant between heterogeneous detectors. This phenomenon implies that the best matching feature for imitation may not be from the corresponding pyramid level, which is an issue that the purely scale-based heuristics fails to address. Thus, we are motivated to conduct feature matching based on their semantics instead of scales. In addition, the object region on nearby feature levels are activated in a similar manner, while the similarity diminishes as the distance between levels becomes greater. This implies that student features from other levels which carry similar semantics should also be involved during imitation to fully

exploit the potential of teacher’s representation power.

To this end, we propose a semantic-guided feature imitation (SGFI) scheme which performs soft matching between features of the student-teacher pair as illustrated in Figure 3. Specifically, given a proposal indexed by i , the teacher’s feature $T_i \in R^{H \times W \times C}$ is extracted from the heuristically assigned pyramid level (which is consistent with its training process). In contrast, the student’s features from all levels are extracted and mapped to the same dimension as T_i using f_{adapt} and is denoted as $S_i \in R^{L \times H \times W \times C}$, where L is the number of pyramid levels. We first project T_i and each $S_i^l (l = 1, \dots, L)$, onto the same embedding space with $f_{embed} : R^{H \times W \times C} \rightarrow R^{C_{key}}$, then the level-wise weights α_i are calculated using the dot product between the embeddings followed by a softmax function, which is used to aggregate $S_i^l (l = 1, \dots, L)$ to obtain S_{agg_i} . The final loss is the mean square error between T_i and S_{agg_i} . The calculation of cross-level imitation loss can be formulated as:

$$\begin{aligned}
 K_{s_i} &= f_{embed}(f_{adapt}(S_i)), \quad K_{t_i} = f_{embed}(T_i) \\
 \alpha_i &= \text{softmax}\left(\frac{K_{s_i} K_{t_i}^T}{\tau}\right) \\
 S_{agg_i} &= \sum_{l=1}^L \alpha_i^l \times f_{adapt}(S_i^l) \\
 L_{feat} &= \frac{1}{N} \sum_{i=1}^N (MSE(S_{agg_i}, T_i)) \quad (1)
 \end{aligned}$$

where f_{adapt} is implemented as a convolutional layer; f_{embed} is implemented as a lightweight network which consists of 2 convs with stride=2, each followed by ReLU; both networks for f_{adapt} and f_{embed} are excluded during inference; N and L are the number of proposals and the number of levels in the feature pyramid, respectively; MSE is the mean square error; τ is a learnable temperature to control the sharpness of softmax logits.

Our proposed SGFI effectively addresses the misalignment between feature levels of student-teacher pairs, which can be easily extended to heterogeneous detector pairs.

4.2. Exploiting Region Relationship with Contrastive KD (CKD)

Feature imitation methods transfer teacher’s knowledge by maximizing the agreement between features of the same region. However, the structural information encoded in the relationship between different regions is ignored, which may also provide guidance to the student. To this end, we push the envelop further by integrating the region relationship into KD for detection. Inspired by the recent work [27], we propose to incorporate the idea of contrastive learning into our KD framework. The objective is to maximize the agreement between the representations of positive pairs while pushing away those of the negative pairs in the given

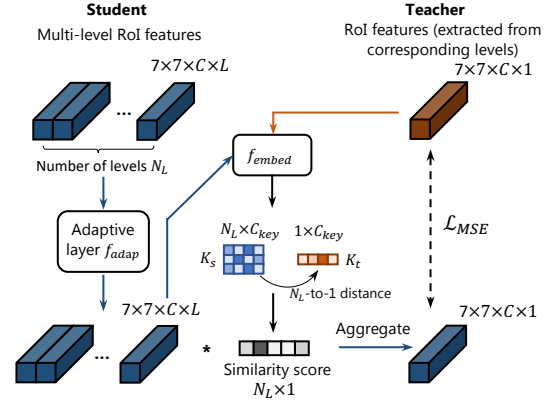


Figure 3. Framework of our semantic-guided feature imitation. Our SGFI automatically performs soft matching between feature pairs across all pyramid levels according to their semantic relations, which is able to provide the optimal guidance to the student.

metric space, which intrinsically captures the relationship between contrastive pairs.

Specifically, given a set B consisting of N RoI bounding boxes, i.e., $B = \{bboxi\}_{i=1, \dots, N}$, their corresponding representations $\{r_s^i, r_t^i\}_{i=1, \dots, N}$ are drawn from the embeddings before the output layer. We form the contrastive pairs as follows: representations that correspond to the same box are constructed as positive pairs while those of different box are constructed as negatives, namely, $x_{pos} = \{r_s^i, r_t^i\}$, $x_{neg} = \{r_s^i, r_t^j\} (i \neq j)$. Our objective is recognizing the positive pair x_{pos} from the set $S = \{x_{pos}, x_{neg}^1, x_{neg}^2, \dots, x_{neg}^K\}$ that contains K negative pairs, which is implemented in the form of InfoNCE loss [19]:

$$L_{ckd} = \frac{1}{N} \sum_{i=1}^N -\log \frac{g(r_s^i, r_t^i)}{\sum_{j=0}^K g(r_s^i, r_t^j)} \quad (2)$$

where K is the number of negative samples; N is the number of proposals in a batch; g is the critic function that estimates the probability of (r_s^i, r_t^i) being the positive pair, which is defined as:

$$g(r_s, r_t) = \exp\left(\frac{f_\theta(r_s) \cdot f_\theta(r_t)}{\|f_\theta(r_s)\| \cdot \|f_\theta(r_t)\|} \cdot \frac{1}{\gamma}\right)$$

where f_θ is a linear function to project the representation to a lower dimension, which is implemented as a fully connected layer whose parameters are shared between the student-teacher pair; $\frac{f_\theta(r_s) \cdot f_\theta(r_t)}{\|f_\theta(r_s)\| \cdot \|f_\theta(r_t)\|}$ is the cosine similarity; γ is a temperature hyper-parameter.

Theoretically, minimizing L_{ckd} is equivalent to maximizing the lower bound of the mutual information between $f_\theta(r_s)$ and $f_\theta(r_t)$ (Detailed proof can be found in previous work [19]):

$$MI(f_\theta(r_s) \cdot f_\theta(r_t)) \geq \log(K) - L_{contrastive}$$

Memory Queue. We implement a memory queue [10] to store the representations for constructing more negative pairs: a queue across multiple GPUs is maintained, and once the max size is reached, the oldest batch is dequeued when new batch arrives. Theoretically, the lower bound becomes tighter as K increases, which implies that using more negative samples benefits representation learning. However, we observe that setting K too large leads to performance degradation. To effect of K is shown in the Appendix.

Negative Sample Assignment Strategy. Another key issue for contrastive KD in detection task is the mechanism to select negative samples. Specifically, the dilemma lies in the overlapping between region proposals: those proposal with large overlaps may contain similar semantics, thus pushing them away may cause instability during training. To address this issue, we use IoU to filter out the highly overlapping proposal boxes and exclude them from negative samples. We conduct an ablative study in the appendix to decide the optimal IoU threshold.

4.3. General Detection KD Framework (G-DetKD)

In some specific scenarios, only detectors with certain architectures can be deployed due to hardware constraints, which may cause the student to have a different architecture from the teacher. Thus, it is promising if knowledge distillation can be conducted between heterogeneous detector pairs. However, previous works only consider KD between homogeneous detectors pairs due to their lack of extensibility. Therefore, we are motivated to propose a general detection KD framework applicable for both homogeneous and heterogeneous student-teacher detector pairs.

4.3.1 Homogeneous Detector Pairs

Homogenous detector pairs are strictly aligned in terms of network categories and feature representations, which facilitates the design of KD framework. Other than the previously introduced SGFI and CKD, we propose two additional techniques to further promote the framework’s effectiveness, namely, **class-aware localization KD** and **head transfer**. In particular, we study the case when the student and teacher are both two-stage detectors.

Class-aware Localization KD. A core component distinguishing detection KD from classification KD lies in how to effectively transfer teacher’s localization ability to student. Intuitively, simply imitating the four coordinates outputted by the teacher provides limited information as they are only “inaccurate targets”. This motivates us to incorporate teacher’s localization knowledge with the class “uncertainty”, i.e., utilizing all the class-wise localization predictions generated by the teacher. For illustration, when the detector captures only a part of an object, how the box should be shifted may depend on what class the object belongs to.




Figure 4. Framework of semantic-guided feature distillation and contrastive distillation for heterogeneous student detectors. As shown in top of the picture, the proposals are sampled from bounding boxes predicted by the localization head, which are then used by SGFI (illustrated in Fig. 3) and CKD; the bottom part illustrates the CKD process: student’s classification features are constructed as contrastive pairs with teacher’s features. In the case of one-stage student, each feature is further decoupled for different box regions centered at the same pixel.

(This idea requires both student and teacher to have localization predictions for each possible class, which is a common setting for two-stage detectors) We calculate the sum of regression values weighted by classification scores:

$$L_{reg} = \frac{1}{N} \sum \left| \sum_{i=0}^C p_t^i \times (reg_t^i - reg_s^i) \right| \quad (3)$$

where C is the number of classes; p_i and reg_i are the classification score and regression outputs of foreground class i ; superscripts s and t represent student and teacher. This class-aware regression loss helps student acquire localization knowledge from teacher; the experiment is shown in Appendix.

Head Transfer. For homogeneous detector pairs, the student and the teacher have backbones with different capacities, while sharing the same head structure. This motivates a more straightforward knowledge transfer strategy for homogeneous detector pairs: directly copying the weights from the teacher’s head to the student. Our experiments demonstrate it helps accelerate students’s convergence and further improves the performance.

Overall, the loss function of our framework for homogeneous detectors can be formulated as: $L = L_{gt} + L_{feat} + L_{ckd} + L_{cls} + L_{reg}$, where L_{gt} is the ground truth loss; L_{feat} , L_{ckd} and L_{reg} correspond to 1, 2 and 3, respectively; L_{cls} is defined in Section 3.

4.3.2 Heterogeneous Detector Pairs

Modern detectors have diversified into multiple families, such as one-stage, two-stage and anchor-free detectors.

Each family has its own merits and weaknesses. In particular, two-stage detectors often have higher performance, while being slower in inference speed. On the other hand, dense prediction detectors (e.g., one-stage and anchor-free detectors) are faster than two-stage counterparts while being less accurate, as they adopt fully convolutional network. In practice, it is a natural idea to use two-stage detectors as teacher to enhance detectors belonging to other families.

The difficulty in knowledge distillation between heterogeneous detector pairs lies in the following aspects: (1) the feature levels are not aligned, e.g., FasterRCNN constructs the feature pyramid from $P2$ to $P6$, while RetinaNet uses $P3$ to $P7$, which creates obstacle for feature distillation; (2) different loss function are used during training, causing their outputs to carry different meanings. E.g., two-stage detectors use cross entropy loss, and dense prediction detectors often adopt focal loss [16], which hinders distillation on the prediction outputs.

We first try to conduct KD on the prediction outputs of FasterRCNN teacher and RetinaNet student and find it bringing only limited gains, and the gains diminish when applied with other KD methods (the results are shown in Table 1). Next, we elaborate on how to extend our feature level distillation method to heterogeneous detectors.

The overview of our G-DetKD for heterogeneous detector pairs is shown in Figure 4. We use the student’s bounding box predictions to generate a set of RoIs, which are then applied to cross-level feature imitation and contrastive KD in the same way as described in Sections 4.1 and 4.2.

RoI Extraction. Analogous to the RPN head in two-stage detectors, we first match all the student’s predicted boxes with the ground truths and label those with IoUs greater than a threshold (set to be 0.5) as positive samples. Then the boxes are sampled with 1:3 ratio of positive to negative samples.

Semantic-guided Feature Imitation. As introduced in 4.1, our method eliminates the requirement for strict alignment between feature levels of the student-teacher pairs, thus is directly applicable to heterogeneous detector pairs.

Decoupling Representations for Contrastive KD. The typical bounding box head structures of dense prediction detector consists of separate multi-convolution classification and localization branches. As illustrated in Figure 4, given a set of RoIs, the contrastive pairs are constructed using the features of the teacher’s last fully connected layer and the corresponding features from the last layer of student’s classification branch. For one-stage detectors, each representation encodes information of multiple anchors centered at the same location, which is decoupled for each anchor.

Overall, the loss function of our framework for heterogeneous detectors can be formulated as: $L = L_{gt} + L_{feat} + L_{ckd}$, where L_{gt} is the ground truth loss; L_{feat} and L_{ckd}

Part	AP
Baseline R18	32.3
Cls pred	32.6 ^{+0.3}
Reg pred	32.7 ^{+0.4}
SGFI	36.1^{+3.8}
SGFI+pred(cls+reg)	36.1 ^{+3.8}

Table 1. Knowledge distillation between R18-RetinaNet student and R50-fasterRCNN teacher. The gain of prediction KD is negligible and diminishes when applied with our SGFI.

correspond to 1,2 and 3, respectively.

5. Experiments

Datasets. We evaluate our knowledge distillation framework on various modern object detectors and popular benchmarks. Our main experiments are conducted on COCO dataset [17]. When compared with other popular algorithms, `test-dev` split is used and the performances are obtained by uploading the results to the COCO test server. **Berkeley Deep Drive** (BDD) [34] and **PASCAL VOC** (VOC) [6] are then used to validate the generalization capability of our method. The default evaluation metrics for each dataset is adopted.

Implementation Detail. We use cosine annealing learning rate schedule; the initial learning rate is set to 0.3. Training is conducted on 8 GPUs using synchronized SGD, batch size is 2 for each GPU. The shorter side of the input image is scaled to 800 pixels, the longer side is limited up to 1333 pixels. ‘1x’ (namely 12 epochs) and ‘2x+ms’ (namely 24 epochs with multi-scale training) training schedules are used. More details can be found in Appendix.

5.1. Main Results

Decoupling the effectiveness of each distillation component. We first investigate the effectiveness of each distillation component for three types of student detectors. The experiment results are shown in Table 2, where use FasterRCNN-R50-FPN as teacher and all students have R18-FPN backbones. The result for FasterRCNN student demonstrates the effectiveness of our class-aware localization KD, which outperforms the naive localization KD approach by a large margin. In addition, contrastive KD is the most effective approach when used alone, which leads to 3.1% gain in *AP*. On the other hand, for heterogeneous students, we can see our SGFI and Contrastive KD provide even higher performance gain compared to their results for homogeneous pairs. The combination achieving the highest performance gain is selected for other main experiments.

Comparison with SOTA Detection KD Methods. We compare our results for homogeneous pairs with previous works in Table 3, since those methods only consider distillation between homogeneous teacher-student pairs. The results show that our G-DetKD and CKD consistently out-

KD Method	AP		
	FasterRCNN[23]	Retina[16]	FCOS[28]
Teacher R50-FPN (2x+ms)	39.9	-	-
Student R18-FPN	34.0	32.6	30.3
SGFI	36.7 ^{+2.7}	36.0 ^{+3.4}	35.2 ^{+4.9}
CKD	†37.1 ^{+3.1}	35.8 ^{+3.2}	34.9 ^{+4.6}
Pred (Cls+Reg)	36.4 ^{+2.4}	33.1 ^{+0.5}	-
Pred (Cls+CAREg)	37.0 ^{+3.0}	-	-
HT	35.3 ^{+1.3}	-	-
SGFI+CKD	37.6 ^{+3.6}	36.3 ^{+3.7}	35.6 ^{+5.3}
SGFI+CKD+Pred+HT	37.9 ^{+3.9}	-	-

Table 2. Effectiveness of each KD component for different type of students. SGFI and CKD are our semantic-guided feature imitation and contrastive KD, respectively; Cls and CAREg are classification KD and our class-aware regression KD, respectively; HT is our head transfer technique.

Method	AP	AP _S	AP _M	AP _L
Teacher R152-FPN	41.5	24.1	45.8	54.0
Student R50-FPN	37.4	21.8	41.0	47.8
FGFI[29]	39.8 ^{+2.4}	22.9	43.6	52.8
TADF[25]	40.0 ^{+2.6}	23.0	43.6	53.0
CKD (Ours)	40.3^{+2.9}	23.2	44.1	52.3
G-DetKD (Ours)	41.0^{+3.6}	23.7	45.0	53.7
Teacher R50-FPN	37.4	21.8	41.0	47.8
Student R50(1/4)-FPN	29.4	16.3	31.7	39.0
FGFI[29]	31.7 ^{+2.3}	17.1	34.2	43.0
CKD (Ours)	32.4^{+3.0}	17.2	34.8	43.0
G-DetKD (Ours)	33.7^{+4.3}	18.1	36.6	44.5

Table 3. Comparison between our KD methods with other approaches. The results show that our G-DetKD and CKD consistently outperforms others by a large margin.

performs others by a large margin.

Exploit the Full Potential of Whole Framework. In order to further explore the full potential of our G-DetKD, we use a more powerful X101-FasterRCNN-InstaBoost model as the teacher and train the student for using 2x+ms schedule. We conduct experiments on students belonging to four families. The results shown in Table 4 demonstrate the significant performance gains brought by our method: +3.8 for R101-RetinaNet, +3.1 for R101-FCOS, +3.4 for R101-FasterRCNN and +4.1 for X101-Cascade (with X101-HTC teacher). In addition, we compare the performance of our KD-enhanced detectors with the current SOTA detector designs on COCO test-dev split in Table 5. The results show that our KD method can upgrade regular detectors to easily dominate the SOTA methods without any modification to the detector. Since KD and detector design are parallel approaches, there is potential for further improvement if we apply our KD framework to the best detector designs.

Student		Teacher	AP
R18 (34.2)	RetinaNet[16]	X101	38.5 ^{+4.3}
R50 (38.5)			43.3 ^{+4.8}
R101 (40.8)			44.6 ^{+3.8}
R18 (34.4)	FCOS[28]	FasterRCNN	38.6 ^{+4.2}
R50 (38.6)			43.1 ^{+4.5}
R101 (41.5)			44.6 ^{+3.1}
R18 (36.0)	FasterRCNN[23]	InstaBoost[7](44.5)	39.9 ^{+3.9}
R50 (39.9)			44.0 ^{+4.1}
R101 (41.8)			45.2 ^{+3.4}
X101-Cascade(46.0)	CascadeRCNN[1]	X101-HTC[3](50.4)	50.1^{+4.1}

Table 4. Our detection KD framework brings significant performance boosts for both homogeneous and heterogeneous detector pairs. All students use two-stage detectors as teachers. “Insta” means InstaBoost [7]; “HTC” stands for Hybrid Task Cascade [3].

Method	backbone	AP	AP _{@.5}	AP _{@.7}
RepPoints [31]	R101-DCN	45.0	66.1	49.0
SAPD [38]	R101	43.5	63.6	46.5
ATSS [37]	R101	43.6	62.1	47.4
PAA [13]	R101	44.8	63.3	48.7
BorderDet [20]	R101	45.4	64.1	48.8
BorderDet [20]	X101-64x4d-DCN	47.2	66.1	51.0
RetinaNet (Ours)	R101	44.8	64.2	48.3
FCOS (Ours)	R101	45.0	64.1	48.5
FasterRCNN (Ours)	R101	45.6	65.9	49.9
Cascade (Ours)	X101-32x4d-DCN	50.5	69.3	55.1

Table 5. Comparison of our G-DetKD with SOTA detector design approaches. Detectors in the four last rows are enhanced by our G-DetKD framework.

5.2. Ablation Study

Comparison of Feature Imitation Strategies. We evaluate different feature imitation strategies, including various mask-based methods, RoI feature imitation (RoIFI) and our SGFI, for both homogeneous and heterogeneous student-teacher pairs in Table 6. For homogeneous detectors, mask-based approaches show similar performances, which are outperformed by RoI-based imitation methods by a large margin. Our proposed SGFI beats the best mask-based approach by 1.1% AP. For heterogeneous detectors, the superiority of SGFI becomes more evident, which indicates that SGFI can well resolve the misalignment issue between feature levels of the student-teacher pairs.

Visualizing Pyramid Level Matching in SGFI. To investigate the feature’s matching pattern in SGFI, we visualize the distributions of the pyramid level difference between the student-teacher pair’s best matching features in Figure 5. Specifically, 500 images (consisting of around 50000 positive proposals) are fed into the trained student-teacher pair. For each proposal, we collect its corresponding teacher’s pyramid level and student’s best matching level to calcu-

Method	AP		
	FasterRCNN	Retina	FCOS
Teacher R50-FPN (2x+ms)	39.9	-	-
Student R18-FPN	34.0	32.6	30.3
Whole Feature [2]	35.2 ^{+1.2}	-	-
Anchor Mask [29]	35.6 ^{+1.6}	-	-
Gaussian Mask [25]	35.4 ^{+1.4}	-	-
GT Mask	35.5 ^{+1.5}	34.3 ^{+1.7}	31.8 ^{+1.5}
RoIFI	36.4 ^{+2.4}	35.5 ^{+2.9}	34.7 ^{+4.4}
SGFI	36.7^{+2.7}	36.0^{+3.4}	35.2^{+4.9}

Table 6. Comparison between different feature distillation techniques. As can be seen, our SGFI outperforms the other methods by a large margin. RoIFI stands for RoI feature imitation method.




Figure 5. Distributions of the pyramid level difference between the student-teacher pair’s best matching features.




Figure 6. Visualization of the difference between the student-teacher pair’s correlation matrices at the classification logits. The intensity of the color represents the magnitude.

late the difference. The distributions show that, while most of the best matching features are from the same level, feature matchings from different levels exist. In addition, the possibility of matching diminishes as features’ corresponding levels are farther apart. Note that the starting index of pyramids in RetinaNet and FCOS is larger than that of the teacher by 1, as explained in 3, thus causing the distribution to be centered at -1.

Visualizing the Influence of Contrastive KD at Class Outputs. To verify the effectiveness of our contrastive KD, we visualize the difference between the correlation matrices of the student-teacher pair’s classification logits. As can be seen in fig 6, in comparison with baseline training and prediction KD strategy, the resulting student model of our contrastive KD approach has higher correlation in the classification outputs with the teacher, which indicates that our contrastive KD effectively captures the inter-class correlations and helps the student learn the complex structural information (i.e., the interdependencies between output dimensions) from the teacher.

Generalization Ability of our KD framework. The generalization ability is important for training strategies, for

Student	Teacher	AP
R18 (34.0)	R50 (39.9)	37.9 ^{+3.9}
R50 (37.4)	R50 (39.9)	39.8 ^{+2.4}
R101 (39.4)	R50 (39.9)	40.4 ^{+1.0}
R50 (37.4)	R18 (36.0)	37.7 ^{+0.3}
R50 (37.4)	R50 (39.9)	39.8 ^{+2.4}
R50 (37.4)	R101 (41.8)	41.1 ^{+3.7}

Table 7. Our KD framework consistently boosts the performance given students and teachers with different capacities. The values in the parentheses indicate baseline APs.

which we explore from two aspects: model and dataset. Homogeneous detector pairs are used throughout the experiments.

Generalization for Students and Teachers. A robust KD framework should generalize for models with different capacities. We explore our G-DetKD’s generalization ability in Table 7. The results show that our method consistently results in performance gains for student-teacher pairs with different capacities. Interestingly, we find that student’s performance is boosted even when the teacher is weaker. We assume this phenomenon is due to the regularization effect of our KD method, which is also observed in [35].

Generalization on other Datasets. We show our method generalizes well for different datasets by experimenting on two additional datasets, namely Pascal VOC [6] and BDD [34]. We do not deliberately pick powerful teachers (R101-FPN are used as teachers and is trained under the same schedule as the student). Our R50-FPN achieves 3.8% gain in $AP_{@.9}$ on VOC and 1.4% AP gain on BDD.

Model	$AP_{@.5}$	$AP_{@.7}$	$AP_{@.9}$
R50-FPN	81.8 [↑] 83.0	67.5 [↑] 69.7	16.8 [↑] 20.6
R101-FPN	83.3 [↑] 83.7	69.7 [↑] 70.3	19.8 [↑] 22.4

(a) VOC

Model	AP	AP_S	AP_M	AP_L
R50-FPN	36.5 [↑] 37.9	14.5 [↑] 14.6	39.0 [↑] 45.0	56.9 [↑] 58.5
R101-FPN	37.5 [↑] 38.0	14.5 [↑] 14.7	40.0 [↑] 40.6	57.7 [↑] 58.4

(b) BDD

Table 8. Generalization on VOC and BDD datasets. The value at the left of \uparrow is the baseline AP while the value at the right is AP achieved with KD.

6. Conclusion

In this paper, we propose a semantic-guided feature imitation method and contrastive distillation for detectors, which helps the student better exploit the learnt knowledge from teacher’s feature pyramid. Furthermore, a general KD framework for detectors is proposed, which is applicable for both homogeneous and heterogeneous detector pairs.

References

- [1] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. 5.1
- [2] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, pages 742–751, 2017. 1, 2, 3, 5.2
- [3] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4974–4983, 2019. 5.1, 4
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2
- [5] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6569–6578, 2019. 2
- [6] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, 2010. 1, 5, 5.2
- [7] Hao-Shu Fang, Jianhua Sun, Runzhong Wang, Minghao Gou, Yong-Lu Li, and Cewu Lu. Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 682–691, 2019. 5.1, 4
- [8] Tommaso Furlanello, Zachary C Lipton, Michael Tschanen, Laurent Itti, and Anima Anandkumar. Born again neural networks. *arXiv preprint arXiv:1805.04770*, 2018. 1, 2
- [9] Golnaz Ghiasi, Tsung-Yi Lin, and Quoc V. Le. Nas-fpn: Learning scalable feature pyramid architecture for object detection. In *CVPR*, 2019. 2
- [10] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 2, 4.2
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1, 2
- [12] Chenhan Jiang, Hang Xu, Wei Zhang, Xiaodan Liang, and Zhenguo Li. Sp-nas: Serial-to-parallel backbone search for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11863–11872, 2020. 2
- [13] Kang Kim and Hee Seok Lee. Probabilistic anchor assignment with iou prediction for object detection. *arXiv preprint arXiv:2007.08103*, 2020. 5.1
- [14] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1, 2
- [15] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 2
- [16] Tsung-Yi Lin, Priyal Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *TPAMI*, 2018. 4.3.2, 5.1, 5.1
- [17] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 1, 5
- [18] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 2
- [19] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 4.2, 4.2
- [20] Han Qiu, Yuchen Ma, Zeming Li, Songtao Liu, and Jian Sun. Borderdet: Border feature for dense object detection. *arXiv preprint arXiv:2007.11056*, 2020. 5.1
- [21] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, 2017. 2
- [22] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 1, 2, 5.1, 5.1
- [24] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 2
- [25] Ruoyu Sun, Fuhui Tang, Xiaopeng Zhang, Hongkai Xiong, and Qi Tian. Distilling object detectors with task adaptive regularization, 2020. 1, 2, 3, 5.1, 5.2
- [26] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017. 1, 2
- [27] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv preprint arXiv:1910.10699*, 2019. 2, 4.2
- [28] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 2, 5.1, 5.1
- [29] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019. 1, 3, 5.1, 5.2
- [30] Hang Xu, Lewei Yao, Wei Zhang, Xiaodan Liang, and Zhenguo Li. Auto-fpn: Automatic network architecture adaptation for object detection beyond classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6649–6658, 2019. 2
- [31] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detec-

- tion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9657–9666, 2019. [2](#), [5.1](#)
- [32] L. Yao, R. Pi, H. Xu, W. Zhang, Z. Li, and T. Zhang. Joint-detnas: Upgrade your detector with nas, pruning and dynamic distillation. In *Computer Vision and Pattern Recognition*, 2021. [2](#)
- [33] Lewei Yao, Hang Xu, Wei Zhang, Xiaodan Liang, and Zhenguo Li. Sm-nas: Structural-to-modular neural architecture search for object detection. *arXiv preprint arXiv:1911.09929*, 2019. [2](#)
- [34] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018. [1](#), [5](#), [5.2](#)
- [35] Li Yuan, Francis EH Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3903–3911, 2020. [5.2](#)
- [36] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *arXiv preprint arXiv:1612.03928*, 2016. [1](#), [2](#)
- [37] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9759–9768, 2020. [5.1](#)
- [38] Chenchen Zhu, Fangyi Chen, Zhiqiang Shen, and Marios Savvides. Soft anchor-point object detection. *arXiv preprint arXiv:1911.12448*, 2019. [5.1](#)