Towards Compact Single Image Super-Resolution via Contrastive Self-distillation

Yanbo Wang¹, Shaohui Lin^{1*}, Yanyun Qu²,
Haiyan Wu¹, Zhizhong Zhang¹, Yuan Xie^{1*} and Angela Yao³

¹East China Normal University

²Xiamen University

³National University of Singapore

{51205901021, 51194501183}@stu.ecnu.edu.cn, {shlin,yxie,zzzhang}@cs.ecnu.edu.cn,
yyqu@xmu.edu.cn, ayao@comp.nus.edu.sg,

Abstract

Convolutional neural networks (CNNs) are highly successful for super-resolution (SR) but often require sophisticated architectures with heavy memory cost and computational overhead, significantly restricts their practical deployments on resource-limited devices. In this paper, we proposed a novel contrastive self-distillation (CSD) framework to simultaneously compress and accelerate various off-the-shelf SR models. In particular, a channel-splitting super-resolution network can first be constructed from a target teacher network as a compact student network. Then, we propose a novel contrastive loss to improve the quality of SR images and PSNR/SSIM via explicit knowledge transfer. Extensive experiments demonstrate that the proposed CSD scheme effectively compresses and accelerates several standard SR models such as EDSR, RCAN and CARN. *Code is available at <https://github.com/Booooooo/CSD>.*

1 Introduction

Single Image Super-Resolution (SISR) aims to reconstruct a high-resolution (HR) image given a low-resolution (LR) image. Recently, convolutional neural networks (CNNs) [Kim *et al.*, 2016; Lim *et al.*, 2017] have dominated SR approaches by directly learning a mapping from LR input to HR output with a deep neural network. However, the existing methods [Zhang *et al.*, 2018; Wang *et al.*, 2018] tend to explore complex and delicate network architectures for recovering edge structures and missing texture details. These networks consume significant amounts of memory and computational cost and are therefore impractical to apply in resource-limited devices such as wearables and IoT.

Several SR model compression methods have been proposed to remove the redundant parameters. Recursive models [Yang *et al.*, 2018; Li *et al.*, 2019] share the SR network’s major blocks and reduces the model size and number of parameters. However, the recursive nature still results in time-consuming inference procedures as the model cost up to 15s

Figure 1: Visual comparison of CSSR-Net with different losses.

per 2K image with $4\times$ SR on one Intel i9-10980XE CPU [Li *et al.*, 2019]. Parameter quantization [Ma *et al.*, 2019; Li *et al.*, 2020] can reduce memory storage by converting the parameters into lower bits. However, the computational complexity is still high as the models are rarely fully quantized. For example, [Ma *et al.*, 2019] quantize parameters but still use full-precision activations. Similarly, [Li *et al.*, 2020] uses additional float scales and low-bit weights/activations to approximate the original full-precision weights/activations, still requiring float-based convolutional computation.

We aim to simultaneously compress and accelerate SR models. We propose a simple *self-distillation* framework inspired by [Yu *et al.*, 2019], in which a student network is split from a teacher (target) network by using parts of teacher’s channels in each layer. We term this student network as *Channel-Splitting Super-Resolution network* (CSSR-Net). The teacher and student networks are trained jointly to form two SR models with different computation. According to different computation resources in devices, we can dynamically allocate these two models, *i.e.* select CSSR-Net if exceeding the required computation overhead in the limited-resource devices, and the teacher model otherwise.

Training CSSR-Net and its teacher network jointly using a reconstruction loss *implicitly* transfers knowledge from teacher to CSSR-Net. Implicit knowledge transfer considers to transfer the knowledge from the teacher to CSSR-Net by weight sharing and joint training, which only provides limited internal knowledge. This will result in low PSNR/SSIM and low quality HR images on CSSR-Net. Recently, percep-

*Corresponding author

tual loss [Johnson *et al.*, 2016] have been widely used as additional knowledge to improve the quality, but the results of CSSR-Net still suffer from blurred lines and edges, as shown in Fig. 1. This is due to limited external knowledge of the perceptual loss that uses only ground-truth HR images as the upper bound of CSSR-Net and teacher. This begs our rethinking: *Why not select multiple negative samples as the lower bounds to reduce the optimization space and provide more explicit knowledge to improve the student’s performance?*

To answer the above question, we propose a *contrastive self-distillation (CSD)* scheme to explicitly transfer the knowledge from teacher to CSSR-Net using contrastive loss (CL), which is inspired by contrastive learning [Oord *et al.*, 2018; Chen *et al.*, 2020; He *et al.*, 2020]. As shown in Fig. 2, SR model compression and acceleration is accomplished by CSSR-Net as a student. Meanwhile, CSD receives rich internal and external knowledge from reconstruction error and CL, respectively. There are two “opposing forces” on CL; One pulls the output of CSSR-Net closer to its teacher, the other one pushes the output of CSSR-Net farther away from the negative images in the latent feature space. CL constrains the output of CSSR-Net into the closed upper and lower bounds, which shows better quality images (see Fig. 1).

Our main contributions could be summarized as follows:

- The proposed CSD scheme as a universal method can simultaneously compress and accelerate various SR networks which is also runtime friendly for practical use.
- Self-distillation is introduced to compress and accelerate SR models, while contrastive loss is proposed to further improve the performance of CSSR-Net by effective knowledge transferring.
- Extensive experiments demonstrate the effectiveness of our CSD scheme. For example, on Urban100, the compressed EDSR+ achieves $4\times$ compression rate and $1.77\times$ speedup, with only a minor loss of 0.13db PSNR and 0.0039 SSIM at the resolution scale of $\times 4$.

2 Related Work

2.1 Single Image Super-Resolution

DCNNs based Super-resolution methods [Lim *et al.*, 2017; Zhang *et al.*, 2018; Ahn *et al.*, 2018] has shown impressive performance for SR in recent years. To improve the visual effect of the reconstructed images, the perceptual loss [Johnson *et al.*, 2016] and the adversarial loss [Wang *et al.*, 2018] were introduced. Both losses are effective as regularizers, but cannot prevent blurry areas as they are limited by the information of the ground truth images as an upper bound. Moreover, they require heavy parameters and computations to improve the quality of HR images, making them impractical to be employed on resource-limited embedded devices.

Several approaches have been proposed to reduce parameters for SR; they can be divided into recursive-based SR [Yang *et al.*, 2018; Li *et al.*, 2019], quantization-based SR [Ma *et al.*, 2019; Li *et al.*, 2020] and compact architecture-based SR [Zhao *et al.*, 2020; Hui *et al.*, 2019]. Recently, AdderSR [Song *et al.*, 2021] utilize adder neural networks to avoid massive energy consumptions while GhostSR [Nie *et al.*,

2021] reduced parameters by generating ghost features. For recursive-based SR and quantization-based SR, they are difficult to accelerate SR models in practical applications. In contrast, Our CSD scheme can simultaneously compress and accelerate various off-the-shelf SR models. Our CSD can also be integrated with compact architecture based SR models, which are orthogonal to our core contribution.

2.2 Knowledge Distillation

Knowledge Distillation (KD) [Hinton *et al.*, 2015] is a teacher-student framework that transfers information from a full teacher network to a compact student network. [Lin *et al.*, 2019] achieve accelerating and compressing via low-rank decomposition with knowledge transfer. Recently, self-distillation has been proposed [Zhang *et al.*, 2019; Ruiz and Verbeek, 2020], where knowledge is distilled within the network itself. [Liu *et al.*, 2020] proposes a stage-based self-distillation method termed as MetaDistiller by transferring the knowledge from deeper stages to earlier stages. However, few self-distillation methods are exploited for SR task. Different from MetaDistiller, our teacher and student have the same number of stages and knowledge is constructed explicitly via the contrastive loss rather than using a label generator to produce soft targets.

In line with our work, [Yu *et al.*, 2019] proposed slimmable neural networks executing different widths on various classification models for dynamic resource adjustment. Training with different widths and switchable batch normalization is a class of implicit knowledge distillation which works well for image classification. However, it will leads to sub-optimal results [Yu and Huang, 2019] when directly applying these methods into SR task. In contrast, we introduce a contrastive loss to explicitly transfer the knowledge from teacher to CSSR-Net via the closed upper and lower bound constraints, which significantly improve the performance of CSSR-Net.

2.3 Contrastive Learning

Contrastive losses are widely used in self-supervised learning [Oord *et al.*, 2018; He *et al.*, 2020; Chen *et al.*, 2020] and aim to pull the anchor close to positive points while pushing it away from negative points in the representation space. Previous works [Chen *et al.*, 2020; Xu *et al.*, 2020] have applied contrastive learning in high-level tasks. [Tian *et al.*, 2020] presents a contrastive distillation framework to capture correlations of structured representations for image classification. Our work on the other hand uses contrastive learning to provide external knowledge with upper and lower bounds for SR tasks. Instead of using two separate networks, we adopt a self-distillation framework, facilitating dynamic resource adjustment by using different channel rates in the teacher model.

Recently, [Park *et al.*, 2020] employed contrastive learning to improve unpaired image-to-image translation quality. [Wu *et al.*, 2021] proposed a contrastive regularization to improve the performance of various SOTA dehazing networks. Inspired by [Wu *et al.*, 2021], we introduce a contrastive loss for the SR task; contrastive learning is unexplored for SR and different from [Park *et al.*, 2020] and [Wu *et al.*, 2021], we also introduce a new way to generate negative samples.

Figure 2: The framework of our proposed CSD. Here we choose EDSR as our backbone for example. The darker green parts of the CSSR-Net (student) are shared from its teacher. The student and teacher separately produce a high-resolution image constrained by the reconstruction loss. The rich knowledge is constructed by contrastive loss and explicitly transferred from teacher to student, where the output O_S is pulled to closer to the O_T and pushed far away from the negative samples in the embedding space of VGG network.

3 Method

Our CSD contains two parts: CSSR-Net and contrastive loss (CL). First, we describe the CSSR-Net. Then, we present our CL to construct the upper and lower bound for CSSR-Net. Finally, the overall loss function of CSD scheme is presented and solved by a new optimization strategy.

3.1 Channel-Splitting Super-Resolution Networks

The channel-splitting super-resolution network (CSSR-Net) is a network constructed by splitting any CNN-based super-resolution network in the channel dimension. The CSSR-Net can be considered as a student network while the original network from which it is derived is a teacher network. The pair can be used to construct a self-distillation framework and implicitly transferring the knowledge from the teacher to the student. CSSR-Net is entangled with the teacher network as it shares a portion of weights from the teacher, as illustrated in the left panel of Fig. 2. The width of the CSSR-Net is controlled by a manually set scale factor r_w . For example, $r_w = 0.5$, corresponds to the CSSR-Net retaining half of the width or number of channels as the teacher in all layers.

Inspired by [Yu *et al.*, 2019], we can jointly minimize the reconstruction loss for both CSSR-Net and the teacher as:

$$L_{Rec} = \sum_i^N L_1(O_S^{(i)}, I_{GT}^{(i)}) + \lambda_T L_1(O_T^{(i)}, I_{GT}^{(i)}), \quad (1)$$

$$= \sum_i^N L_1(f^S(I^{(i)}, \theta_S), I_{GT}^{(i)}) + \lambda_T L_1(g^T(I^{(i)}, \theta_T), I_{GT}^{(i)}),$$

where $O_S^{(i)} = f^S(I^{(i)}, \theta_S)$ and $O_T^{(i)} = g^T(I^{(i)}, \theta_T)$ are the output of the CSSR-Net f^S and the teacher network g^T on the LR input $I^{(i)}$ with parameters θ_S and θ_T respectively. θ_S is shared from θ_T and satisfy $\theta_S \subset \theta_T$. $I_{GT}^{(i)}$ is the ground-truth HR image and N is the number of training images.

Directly minimizing Eq. 1 via stochastic gradient descent (SGD) leads to sub-optimal results in SR task (See the re-

sults in Table 1). As a result, CSSR-Net and the teacher network converge with worse results than the corresponding individual training. Moreover, the generated SR images have blurry parts, *e.g.* the area on dense straight lines (see Fig. 1). We speculate the following reason: implicit KD is not strong enough to provide insightful information by two independent loss terms. We expect to add explicit knowledge, which can provide richer internal and external knowledge. Therefore, we introduce contrastive learning to explicitly construct a relationship between the student and the teacher, also providing the closed upper and lower bound to improve performances of both CSSR-Net and the teacher. Upper bound is constructed to pull the output of CSSR-Net to teacher's, while lower bound is to constrain CSSR-Net's output to be far away from negative samples (*e.g.* bicubic upsampled images).

3.2 Contrastive Loss

Contrastive learning has improved representation learning by pulling anchors close to positive samples while pushing away negative samples [Oord *et al.*, 2018; Park *et al.*, 2020; Chen *et al.*, 2020; He *et al.*, 2020]. We propose a novel contrastive loss (CL) to explicitly represent knowledge for training both the teacher network and CSSR-Net. For contrastive learning, we need to consider two aspects: One is to construct the "positive" and "negative" samples, the other is to find the latent feature space to compare the samples.

For the former, we construct the output $O_S^{(i)}$ of CSSR-Net and the output $O_T^{(i)}$ of its teacher as the anchor and the positive sample, respectively. More negative samples may better cover the undesired distribution. We thus sample K images (other than the anchor) from the same mini-batch to $I^{(i)}$ as negative samples. Then upsample them to the same resolution as $O_S^{(i)}$ via bicubic interpolation. Each of them is denoted by $O_{Neg}^{(k)}$, $k = 1, 2, \dots, K$. More detailed experimental results on numbers of negative samples are shown in experiments.

For the latent feature space, we use intermediate features

of a pre-trained model ϕ (e.g. VGG [Simonyan and Zisserman, 2014]). Given positive and negative samples, we can construct the contrastive loss as:

$$L_{CL} = \sum_i^N \sum_j^M \lambda_j \frac{d(\phi_j(O_S^{(i)}), \phi_j(O_T^{(i)}))}{\sum_k^K d(\phi_j(O_S^{(i)}), \phi_j(O_{Neg}^{(k)}))}, \quad (2)$$

where $\phi_j, j = 1, 2, \dots, M$ is the intermediate features from the j -th layer of pre-trained model. M is the number of total hidden layers. $d(x, y)$ is the L1-distance loss between x and y . λ_j is the balancing weight for each layer. We do not update the parameters of pre-trained model ϕ during training. The contrastive loss is presented in Fig. 2 right panel; the loss introduces *opposing forces* pulling the output of CSSR-Net $O_S^{(i)}$ to the output of its teacher $O_T^{(i)}$ and pushing $O_S^{(i)}$ to the negative samples $O_{Neg}^{(k)}$. Note that our contrastive loss is different from InfoNCE [Oord *et al.*, 2018], which uses a dot product-based similarity. Instead of this similarity, our L1-distance loss achieves better performance (See experiments for more details). Additionally related to our CL is the perceptual loss [Johnson *et al.*, 2016], which minimizes the distance loss between the student and the ground-truth from multi-layer features of the pre-trained VGG. This is an upper bound to constrain the student network. Unlike the perceptual loss, however, we also adopt multiple negative samples as a lower bound to reduce the solution space and further improve the performance of CSSR-Net and its teacher.

3.3 The Overall Loss and Its Solver

Overall Loss. The overall loss of our CSD scheme is constructed by leveraging contrastive loss Eq. 2 into reconstruction loss Eq. 1, which can be formulated as:

$$L(\theta_S, \theta_T) = L_{Rec} + \lambda_C L_{CL}, \quad (3)$$

where λ_C is a hyper-parameter for balancing L_{Rec} and L_{CL} .

Solver. Since our CSSR-Net and its teacher are entangled, we need to update both gradients from them. One naive solver is to update all parameters (*i.e.*, θ_S and θ_T) by directly minimize Eq. 3 based on SGD. However, the teacher simply becomes weaker rather than maintain good performance.

As a solution, we detach teacher's gradients from the contrastive loss and only update gradients from reconstruction loss. For student, we take normal gradient updating of Eq. 3. Pseudocode of CSD scheme is summarized in Algorithm 1.

4 Experiments

4.1 Experimental Setups

Implementation Details. Our CSD scheme is implemented by PyTorch 1.2.0 and MindSpore 1.2.0 [Huawei, 2020] with one NVIDIA TITAN RTX GPU. The models are trained with ADAM optimizer by setting $\beta_1 = 0.9, \beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The batch size and total epochs are set to 16 and 300 epochs, respectively. The initial learning rate is 10^{-4} and decayed by $10\times$ at every 2×10^5 iterations. For the latent features in Eq. 2, We extract the features from the 1st, 3rd, 5th, 9th and 13th layers of the pre-trained VGG-19 while with the corresponding coefficients $\lambda_i, i = 1, \dots, 5$ to $\frac{1}{32}, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}$ and

Algorithm 1 Pseudocode of CSD in a PyTorch-like style

```
# f_t: Pre-trained teacher network
# width_mult: the width of network f_s
# lt, lc: lambda_t, lambda_c in Eq. 1, Eq. 3
initialize()
for lr, hr in loader: # load a minibatch
    neg = bic(generate()) # negative samples
    o_s = f_s.forward(lr) # anchor
    o_t = f_t.forward(lr) # positive samples
    # reconstruction loss, Eq. 1
    loss = L1(o_s, hr) + lt * L1(o_t, hr)
    # No gradient to o_t from contrastive loss
    o_t.detach()
    # contrastive loss, Eq. 2
    vgg_s, vgg_t, vgg_n = VGG19(o_s, o_t, neg)
    loss += lc * CL(vgg_s, vgg_t, vgg_n)
    loss.backward() # update
```

1, respectively. We set the hyper-parameters λ_T and λ_C in Eq. 2 to 1 and 200, respectively. The scale factor for the width of CSSR-Net is default set to 0.25 if not specially stated. The input is randomly cropped into patches and augmented with random horizontal flip and 90° rotation to produce the fixed 192×192 HR patches during training.

Datasets. We train all SR models with 800 training images on DIV2K and evaluate on the 100 validation images. We additionally test on four SR benchmarks: Set5 [Bevilacqua *et al.*, 2012], Set14 [Zeyde *et al.*, 2010], BSD100 [Martin *et al.*, 2001] and Urban100 [Huang *et al.*, 2015].

Evaluation Metric. We calculate PSNR and SSIM on the Y channel, also evaluate the compression rate and real speedup on one NVIDIA TITANX RTX GPU.

Teacher Backbones. To validate our CSD, we choose models with different sizes and structures (EDSR+ [Lim *et al.*, 2017], RCAN+ [Zhang *et al.*, 2018] and CARN+ [Ahn *et al.*, 2018]) as teacher backbones, upon which we construct their corresponding CSSR-Nets as students. Note that CARN+ is the most compact network with only 1.1M parameters, compared to RCAN+ (15.6M) and EDSR+ (43.1M).

Baselines. We select individually trained CSSR-Net and its teacher as our baseline. We also compare the proposed CSD with CSSR-Net, which only optimizes Eq. 1 via joint training, denoted as J-T1.

4.2 Quantitative and Qualitative Results

Quantitative Results

We first compress EDSR+, as shown in Table 1. EDSR+ 1.0 $\times$ denotes the original EDSR+ model as a teacher, while EDSR+ 0.25 $\times$ presents a student with 2.7M parameters by setting $r_w = 0.25$. Compared to the teacher, EDSR+ 0.25 $\times$ achieves 16.0 $\times$ compression rate. We can observe that (1) Joint training (J-T1) achieves the consistent PSNR/SSIM with baseline (*i.e.* individual training) both in student and teacher at the same SR scale. For dynamic inference, the baseline needs to store two individual models with different

“+” means self-ensemble technique [Lim *et al.*, 2017]) is applied into the networks

 Figure 3: Qualitative Comparison on EDSR+ with $\times 4$ SR. “B-” and CSSR-Net are baseline and EDSR+ $0.25\times$ by J-T1, respectively.

Model	Method	Scale	DIV2K		Set5		Set14		BSD100		Urban100	
			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
EDSR+ $0.25\times(S)$	Baseline	$\times 2$	34.75	0.9463	38.07	0.9607	33.68	0.9179	32.23	0.9004	32.28	0.9298
		$\times 4$	29.00	0.8379	32.23	0.8957	28.63	0.7825	27.60	0.7366	26.07	0.7851
	J-T1	$\times 2$	34.73	0.9461	38.09	0.9608	33.67	0.9179	32.22	0.9002	32.26	0.9296
		$\times 4$	29.00	0.8379	32.22	0.8958	28.63	0.7827	27.60	0.7369	26.06	0.7852
	CSD	$\times 2$	34.85	0.9473	38.19	0.9612	33.85	0.9199	32.29	0.9012	32.55	0.9327
		$\times 4$	29.13	0.8416	32.34	0.8974	28.72	0.7856	27.68	0.7396	26.34	0.7948
EDSR+ $1.0\times(T)$	Baseline	$\times 2$	35.12	0.9699	38.20	0.9606	34.02	0.9204	32.37	0.9018	33.10	0.9363
		$\times 4$	29.33	0.8456	32.60	0.8998	28.90	0.7892	27.76	0.7428	26.76	0.8061
	J-T1	$\times 2$	35.11	0.9489	38.29	0.9615	34.10	0.9216	32.40	0.9024	32.13	0.9371
		$\times 4$	29.32	0.8453	32.59	0.8997	28.89	0.7889	27.76	0.7424	26.73	0.8055
	CSD	$\times 2$	35.15	0.9402	38.31	0.9617	34.07	0.9216	32.41	0.9026	33.17	0.9375
		$\times 4$	29.34	0.8458	32.60	0.9000	28.91	0.7893	27.77	0.7430	26.79	0.8063

 Table 1: Results on EDSR+. S and T respectively indicate EDSR+ $0.25\times$ and EDSR+ $1.0\times$, where the number 0.25 or 1.0 is the width scale. EDSR+ $0.25\times$ has 2.7M parameters, which achieves $16\times$ compression rate compared to T .

parameters, while J-T1 only requires the parameter of the teacher; (2) Compared to baseline and J-T1, CSD achieves the best performance both in student and teacher at the same SR scale, except the SSIM metric of teacher at SR scale of 2 on DIV2K. For example, on Urban100, EDSR+ $0.25\times$ using CSD achieves the PSNR gains over baseline by 0.26dB and 0.24dB at the SR scale of 2 and 4, respectively; (3) Using CSD, our EDSR+ $0.25\times$ model achieves $16\times$ compression rate, only with the loss of 0.2db PSNR at the SR scale of 4 (*i.e.* 29.13dB *vs.* 29.33dB in baseline) on DIV2K.

We further compress RCAN+ and CARN+, which are more compact. Additionally, we set $r_w = 0.5$. As shown in Fig. 4, our CSD achieves higher PSNR at the SR scale of 4 on Urban100, compared to baseline. For example, on Urban100 $\times 4$, our CSD-RCAN+ $0.5\times$ achieves 0.12dB PSNR gains over B-RCAN+ $0.5\times$. We also found that CARN+ compression is relatively difficult, which is due to the smallest number of redundant parameters. However, our CSD still achieves higher PSNR, compared to the baseline based on individual training. Moreover, compared to SOTA SR quantization methods (*i.e.* PAMS-4bit [Li *et al.*, 2020] and PACT-4bit [Choi *et al.*, 2018]), our CSD achieves 0.19dB and 0.37dB PSNR gains over PAMS and PACT, while with the smallest parameter number of 0.3M (*vs.* 0.48M both in PAMS and

PACT). For fair comparison to evaluate the speedup, we report the total inference time across Urban100 dataset with SR scalar of 4 (See Fig. 4(b)). Our CSD achieves $3.9\times$ compression rate and $1.38\times$ GPU speedup rate only with 0.2dB PSNR loss, compared to the original model.

Qualitative Results

As shown in Fig. 3, we compare our CSD with other training methods on the quality of the enhanced images at $4\times$ SR scale on Urban100, BSD100 and DIV2K. For simplicity, we select the student model EDSR+ $0.25\times$ for visual comparison. We can see that both individually trained (baseline) and jointly trained (CSSR-Net) output blurry images, especially in the areas where straight lines are very dense. In contrast, our CSD can alleviate the blurry problem, which means that high frequency information is effectively recovered.

4.3 Ablation Study

To demonstrate the effectiveness of the proposed CSD scheme, we conduct ablation study to analyze the effect of contrastive loss, S-T distillation, updating strategies and the number of negative samples. EDSR+ is selected as our backbone in ablation study.

Effect of InfoNCE Loss, Perceptual Loss and CL. InfoNCE loss [Oord *et al.*, 2018; He *et al.*, 2020] has been widely used in contrastive learning. It uses the dot-product

This number has been converted to 32-bit float number.

(a) PSNR-parameter

(b) PSNR-Speed

 Figure 4: Compression results on Urban100 $\times 4$, GPU speedup rate and PSNR of their corresponding students (width=0.5).

Methods	DIV2K		Urban100	
	PSNR	SSIM	PSNR	SSIM
Baseline	29.00	0.8379	26.07	0.7851
CSD (InfoNCE)	29.10	0.8402	26.24	0.7907
J-T (W/ Perceptual Loss)	28.72	0.8317	25.72	0.7757
W/O T (GT Pos.)	29.11	0.8406	26.30	0.7924
CSD (GT Pos.)	28.87	0.8356	26.00	0.7844
T-S Separate	29.03	0.8406	26.15	0.7894
CSD (Ours)	29.13	0.8416	26.34	0.7948

 Table 2: Comparison of different contrastive losses and S-T distillations on EDSR+ 0.25 $\times$ model on Urban100 $\times 4$.

operation to measure the distance between two vectors. Instead of InfoNCE loss, we use contrastive loss in Eq. 2 based on L1-distance measurement. As shown in Tab. 2, our CL achieves the best results, compared to InfoNCE and J-T with perceptual loss. For example, Our CSD based on L1-distance achieves 0.0014 and 0.0041 SSIM gains over InfoNCE based CSD on DIV2K and Urban100, respectively.

Effect of S-T Distillation We further evaluate the effectiveness of our CSSR-Net. We consider three S-T distillation strategies: (1) W/O T (GT Pos.) only use ground-truth HR images as positive samples, which remove the teacher’s branch; (2) CSD (GT Pos.) uses ground-truth HR images as positive samples in CSD scheme; (3) T-S Separate where the student does not share weights with the teacher.

As shown in Table 2, our CSSR-Net enables self-distillation, which receives richer information from its entangled teacher. It is worth noticing that the performance of CSD is better than CSD (GT Pos.). We speculate that HR images provide a stronger upper bound which is more difficult for the limited capacity S to fully exploit. This is consistent with findings in [Cho and Hariharan, 2019], where better teachers do not necessarily help to train better students.

Effect of Updating Strategy. Fig. 5(a) presents a comparison of different updating strategies on CL. CSD-A re-

 Figure 5: Comparison of updating strategies and numbers of negative samples on Urban100 $\times 4$

moves the reconstruction loss for T in Eq. 3; CSD-B directly minimize Eq. 3 by both updating weights of S and T from the gradient of CL; CSD detaches T ’s gradient from CL. Obviously, our CSD achieves significantly higher performance, compared to CSD-A and CSD-B. In CSD-A, T collapses quickly in the early stage during training and is unable to provide valid information to S , resulting in much worse performance. In CSD-B, T tends to learn the student such that T become weaker during training.

Effect of the Number of Negative Samples. Finally, we explore the effect of the number of negative samples. As shown in Fig. 5(b), adding more negative samples achieves better performance. However, training memory and time grows as the number of negative samples increases. For the performance-efficiency trade-off, we choose the number of 10 negative samples in our all experiments.

5 Conclusion

In this paper, we propose a novel CSD scheme for simultaneously compressing and accelerating SR models, which consists of CSSR-Net and contrastive loss. Constructed from the target (teacher) network, CSSR-Net shares the part weights, directly compressing the target network. To mine for rich knowledge from the teacher, a novel contrastive loss is proposed for explicit knowledge transfer, which ensures that the output of CSSR-Net is pulled closer to the teacher’s and pushed far away from the blurry images. We have comprehensively evaluated the performance of CSD scheme for compressing and accelerating various SR models on standard benchmark datasets with superior performance. We believe the proposed CSD scheme can be generalized to other low-level vision tasks (*e.g.* dehazing, denoising and deblurring), which will be explored in future work.

Acknowledgements

This work is sponsored by National Natural Science Foundation of China (61972157, 61772524, 61876161, 61902129), Natural Science Foundation of Shanghai (20ZR1417700), CAAI-Huawei MindSpore Open Fund, National Key Research and Development Program of China (2019YFC1521104), Zhejiang Lab (2019KD0AC02, 2020NB0AB01), the Fundamental Research Funds for the Central Universities, Shanghai Sailing Program (21YF1411200). We thank MindSpore[Huawei, 2020] for their partial support.

References

- [Ahn *et al.*, 2018] Namhyuk Ahn, Byungkun Kang, and Kyung-Ah Sohn. Fast, accurate, and lightweight super-resolution with cascading residual network. In *ECCV*, 2018.
- [Bevilacqua *et al.*, 2012] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *BMVC*, 2012.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [Cho and Hariharan, 2019] Jang Hyun Cho and Bharath Hariharan. On the efficacy of knowledge distillation. In *ICCV*, 2019.
- [Choi *et al.*, 2018] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. In *CVPRW*, 2018.
- [He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020.
- [Hinton *et al.*, 2015] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [Huang *et al.*, 2015] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, 2015.
- [Huawei, 2020] Huawei. Mindspore. <http://www.mindspore.cn/>, 2020.
- [Hui *et al.*, 2019] Zheng Hui, Xinbo Gao, Yunchu Yang, and Xiumei Wang. Lightweight image super-resolution with information multi-distillation network. In *ACM-MM*, 2019.
- [Johnson *et al.*, 2016] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [Kim *et al.*, 2016] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Deeply-recursive convolutional network for image super-resolution. In *CVPR*, 2016.
- [Li *et al.*, 2019] Zhen Li, Jinglei Yang, Zheng Liu, Xiaomin Yang, Gwanggil Jeon, and Wei Wu. Feedback network for image super-resolution. In *CVPR*, 2019.
- [Li *et al.*, 2020] Huixia Li, Chenqian Yan, Shaohui Lin, Xiawu Zheng, Yuchao Li, Baochang Zhang, Fan Yang, and Rongrong Ji. Pams: Quantized super-resolution via parameterized max scale. *arXiv preprint arXiv:2011.04212*, 2020.
- [Lim *et al.*, 2017] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, 2017.
- [Lin *et al.*, 2019] Shaohui Lin, Rongrong Ji, Chao Chen, Dacheng Tao, and Jiebo Luo. Holistic cnn compression via low-rank decomposition with knowledge transfer. *TPAMI*, 41(12):2889–2905, 2019.
- [Liu *et al.*, 2020] Benlin Liu, Yongming Rao, Jiwen Lu, Jie Zhou, and Cho jui Hsieh. Metadistiller: Network self-boosting via meta-learned top-down distillation. In *ECCV*, 2020.
- [Ma *et al.*, 2019] Yinglan Ma, Hongyu Xiong, Zhe Hu, and Lizhuang Ma. Efficient super resolution using binarized neural network. In *CVPRW*, 2019.
- [Martin *et al.*, 2001] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001.
- [Nie *et al.*, 2021] Ying Nie, Kai Han, Zhenhua Liu, An Xiao, Yiping Deng, Chunjing Xu, and Yunhe Wang. Ghostsr: Learning ghost features for efficient image super-resolution. In *CVPR*, 2021.
- [Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [Park *et al.*, 2020] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, 2020.
- [Ruiz and Verbeek, 2020] Adrià Ruiz and Jakob Verbeek. Distilled hierarchical neural ensembles with adaptive inference cost. *arXiv preprint arXiv:2003.01474*, 2020.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Song *et al.*, 2021] Dehua Song, Yunhe Wang, Hanting Chen, Chang Xu, Chunjing Xu, and Dacheng Tao. Addersr: Towards energy efficient image super-resolution. In *CVPR*, 2021.
- [Tian *et al.*, 2020] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020.
- [Wang *et al.*, 2018] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *ECCV*, 2018.
- [Wu *et al.*, 2021] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Contrastive learning for compact single image dehazing. In *CVPR*, 2021.
- [Xu *et al.*, 2020] Guodong Xu, Ziwei Liu, Xiaoxiao Li, and Chen Change Loy. Knowledge distillation meets self-supervision. In *ECCV*, 2020.
- [Yang *et al.*, 2018] Xin Yang, Haiyang Mei, Jiqing Zhang, Ke Xu, Baocai Yin, Qiang Zhang, and Xiaopeng Wei. Drfn: Deep recurrent fusion network for single-image super-resolution with large factors. *TMM*, 21(2):328–337, 2018.
- [Yu and Huang, 2019] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *ICCV*, 2019.
- [Yu *et al.*, 2019] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. In *ICLR*, 2019.
- [Zeyde *et al.*, 2010] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Int. Conf. Curves Surf*, 2010.
- [Zhang *et al.*, 2018] Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-resolution using very deep residual channel attention networks. In *ECCV*, 2018.
- [Zhang *et al.*, 2019] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *ICCV*, 2019.
- [Zhao *et al.*, 2020] Hengyuan Zhao, Xiangtao Kong, Jingwen He, Yu Qiao, and Chao Dong. Efficient image super-resolution using pixel attention. In *ECCVW*, 2020.