You Only Need One Step: Fast Super-Resolution with Stable Diffusion via Scale Distillation

Mehdi Noroozi, Isma Hadji, Brais Martinez, Adrian Bulat, and Georgios Tzimiropoulos

> Samsung AI Cambridge {m.noroozi,isma.hadji}@samsung.com

Abstract. In this paper, we introduce YONOS-SR, a novel stable diffusion based approach for image super-resolution that yields state-ofthe-art results using only a single DDIM step. Specifically, we propose a novel scale distillation approach to train our SR model. Instead of directly training our SR model on the scale factor of interest, we start by training a teacher model on a smaller magnification scale, thereby making the SR problem simpler for the teacher. We then train a student model for a higher magnification scale, using the predictions of the teacher as a target during the training. This process is repeated iteratively until we reach the target scale factor of the final model. The rationale behind our scale distillation is that the teacher aids the student diffusion model training by i) providing a target adapted to the current noise level rather than using the same target coming from ground truth data for all noise levels and ii) providing an accurate target as the teacher has a simpler task to solve. We empirically show that the distilled model significantly outperforms the model trained for high scales directly, especially with few steps during inference. Having a strong diffusion model that requires only one step allows us to freeze the U-Net and fine-tune the decoder on top of it. We show that the combination of spatially distilled U-Net and fine-tuned decoder outperforms state-of-the-art methods requiring 200 steps with only one single step.¹

1 Introduction

Diffusion models have shown impressive performance in various image generation tasks [22, 42], including image super-resolution (SR) [3, 24, 25, 32]. However, the large number of sequential denoising passes required by the sampling strategy results in extreme computational cost, even for stable diffusion-based models (SD) that operate in the latent space of an autoencoder. Recently, several approaches have been proposed to reduce the number of sampling steps [18, 26, 28, 29]. Unfortunately, such approaches usually compromise performance, especially for the lower number of steps.

¹ The code will be available here once all approvals are processed: https://github.com/SamsungLabs/yonos



Fig. 1: Qualitative comparison for $\times 4$ and $\times 8$ magnifications. Each column shows top to bottom LR input image, 1 and 200 step SD-SR, 1-step YONOS-SR(ours). SD-SR represents the standard Stable Diffusion-based SR model. The 1-step SD-SR method lacks quality in terms of detailed textures compared to 200-steps of the same model; see building texture in the first column and hairs in the middle column. In contrast, our method outperforms 200-steps SD-SR with only one step, especially for $\times 8$ magnification where SD-SR fails to recover the details even with 200 steps. Samples are taken from DIV2K validation set. Images are best seen in a display and zoomed in.

Typically, diffusion-based models yield the best results on image patches of similar sizes to those seen during training (e.g. 64×64 for SD [22]). On the other hand, super-resolution applications require operating in high-resolution settings, drastically exacerbating the computational issues of diffusion-based models. For example, a SR model that aims for a magnification of $\times 4$ going from 256×256 to 1024×1024 requires dividing the input image into 16 patches of 64×64 and running the model on each patch individually, making a large number of steps prohibitive for realistic use cases. Using state-of-the-art step-reduction strategy, such as more efficient samplers [18, 19, 28] can partially alleviate this issue but still falls widely short of practical needs. For example, going down to the target of 1 DDIM step results in a significant drop in performance compared to a typical model that does 200 inference steps, as shown in Fig. 1.

One differentiating characteristic of the super-resolution task is that it is conditioned on the low-resolution (LR) input image to yield the target highresolution (HR) image. Unlike the task of text-to-image generation, which relies on text conditioning, the LR image provides closer content to the target HR image, especially at lower scale factors. Therefore, conditioning the diffusion model on the LR image at low-scale factors makes the task inherently simpler for the diffusion model. In this paper, we take advantage of this peculiarity and introduce a novel training strategy dubbed *scale distillation*. While typical diffusion-based SR methods train the model for super-resolution by conditioning directly on the LR image at the target scale factor, we instead propose a progressive training approach, where we start by training a model for lower scale factors (*i.e.* where the conditioning signal is closer to the target) and progressively increase to the target scale factor using the previously trained model as a teacher.

More specifically, instead of using the raw data to train a model for large scale factors, scale distillation obtains a rich and accurate supervisory signal from a teacher trained for a smaller scale factor. We first train a teacher that takes a less degraded image as input and, therefore, has an easier task to solve during training. Then, we train a model for a larger scale factor as a student while initializing it with the same weights as the teacher, which is now frozen. For a given time step during the training, we feed both teacher and student with the same noisy version of the HR image. However, we condition the teacher with the less degraded LR image (*i.e.* using the same scale that was used during teacher training), while we condition the student on the target (more degraded) LR image. We then use the teacher's prediction as a target to train the student.

This training strategy has two direct advantages: i) Unlike typical training where the supervisory signal is somewhat ambiguous as the target is the same for all noise levels, our student receives its target from the teacher and is therefore adaptive to the noise level. ii) The target is more accurate, especially in terms of the finer detail, because the teacher takes a less degraded LR image as input.

The proposed scale distillation approach allows the model to solve the SR task in fewer steps as we have simplified the task for the student. In fact, we show that models trained with our approach improve significantly when a few steps are used during the inference, e.g. one step, see Fig. 3. Therefore, a direct

4 M. Noroozi et al.

advantage of the proposed approach is that fine-tuning the decoder directly on top of the diffusion model becomes computationally tractable due to the single inference step required. Taking advantage of this fine-tuning, we show that You Only Need One Step (YONOS)-SR outperforms state-of-the-art diffusion-based SR methods that require a large number (*e.g.* 200) of inference steps.

In summary, our contributions are threefold: **I**) We introduce scale distillation to train SD models with a more accurate and fine supervisory signal for image super-resolution tasks. **II**) We show that our proposed scale distillation strategy yields more efficient SD models that allow for directly fine-tuning the decoder on top of a frozen one-step diffusion model. **III**) We show that combining scale distillation followed by decoder fine-tuning yields state-of-the-art results on the SR task, even at high magnification factors, while requiring only one step.

2 Related work

Real image super-resolution. Image super-resolution entails restoring a High Resolution (HR) image given its Low Resolution (LR) observation. Solving this task for real images is especially challenging given the dramatic differences in real-world image distributions [10, 11, 17, 38]. These differences arise from varying image degradation processes, different imaging devices, and image signal processing methods, all of which are difficult to properly model and generalize. For this reason, real image super-resolution (or blind super-resolution) has received significant interest among the research community [11,16,32–34,37,38,41]. While some methods attempt to learn the degradation process [5,20,31,39], their success remains limited due to the lack of proper large scale training data [17], even while using some unsupervised methods [44]. In contrast, more popular approaches tackle the problem by explicitly modeling the degradation pipeline to create synthetic LR-HR pairs to use for training [15,27,34,41]. Given, the wider success of the explicit degradation modeling approach, we elect to rely on the widely used RealESRGAN degradation pipeline [34] in training our model.

Diffusion-based super-resolution. Since the early SRCNN [4] method, many deep learning-based solutions for blind super-resolution have been proposed [2, 11, 22, 24, 25, 34, 37, 41, 44]. Early work took advantage of this space by using semantic segmentation probability maps for guiding SR [35]. Most recent methods aim at taking advantage of learned generative priors to simplify the inverse imaging problem of blind image super-resolution. Usually, methods following this paradigm [34, 37, 41] rely on GANs [6]. More recently, diffusion models showed remarkable generative capabilities yielding impressive results across a range of applications [22, 42]. As such, in this paper, we follow several recent works [22, 24, 25, 32] and rely on diffusion-based generative models to tackle the super-resolution problem. While diffusion-based models achieve impressive results, their main shortcoming is the long inference time. Diffusion-based models require several inference steps through the model to yield a final output, thereby limiting their practical use. Therefore, in this paper, we tackle the important

problem of speeding up the inference of diffusion-based super-resolution.

Guided distillation. Recognizing the inference speed shortcoming of diffusion models, several works have been proposed recently to address this issue [18, 19, 21, 26, 28]. These methods can be categorized into two main tacks. One approach tackles this problem at inference time by either proposing more efficient samplers [12, 28] or relying on higher-order solvers [18, 19]. More closely related to ours are methods that aim at directly training a diffusion model that can solve the generative problem at hand in fewer steps through *temporal* distillation [21, 26, 29]. Our method tackles the problem at training time as well but we propose *scale* distillation. Our main idea is to reduce the inference speed by progressively making the generative problem easier during training. Notably, our approach is orthogonal to temporal distillation and can be used in tandem with it.

3 YONOS-SR

In this section, we describe YONOS-SR, our diffusion-based model for image super-resolution. First, we present an overview of the image super-resolution framework with the latent diffusion models in Sec. 3.1. We then discuss our proposed scale distillation method that allows us to improve the performance with fewer sampling steps, *e.g.* 1-step, in Sec. 3.2. Finally, in Sec. 3.3, we discuss how the 1-step diffusion model allows for fine-tuning a decoder directly on top of the diffusion model, with a frozen U-Net.

3.1 Super-resolution with latent diffusion models

Given a training set in the form of pairs of low and high-resolution images $(\mathbf{x}_h, \mathbf{x}_l) \sim p(\mathbf{x}_h, \mathbf{x}_l)$, the task of image super-resolution involves estimating the probability distribution of $p(\mathbf{x}_h | \mathbf{x}_l)$. The stable diffusion framework uses a probabilistic diffusion model applied on the latent space of a pre-trained and frozen autoendoer. Let us assume that $\mathbf{z}_h = \mathcal{E}(\mathbf{x}_h), \mathbf{z}_l = \mathcal{E}(\mathbf{x}_l)$ be the corresponding projection of a given low and high-resolution images $(\mathbf{x}_h, \mathbf{x}_l)$, where \mathcal{E} is the pre-trained encoder. The forward process of the diffusion model, $q(\mathbf{z} | \mathbf{z}_h)$ is a Markovian Gaussian process defined as

$$q(\mathbf{z}_t | \mathbf{z}_h) = \mathcal{N}(\mathbf{z}_t; \alpha_t \mathbf{z}_h, \sigma_t \mathbf{I}), \ \mathbf{z} = \{\mathbf{z}_t | t \in [0, 1]\}$$
(1)

where \mathbf{z} denotes the latent variable of the diffusion model and α_t, σ_t define the noise schedule such that the log signal-to-noise ratio, $\lambda_t = \log[\alpha_t^2/\sigma_t^2]$, decreases with t monotonically. During training, the model learns to reverse this diffusion process progressively, *i.e.* estimate $p(\mathbf{z}_{t-1}|\mathbf{z}_t)$, to generate new data from noise.

The super-resolution objective function is derived by maximizing a variational lower bound of the data log-likelihood of $p(\mathbf{z}_h | \mathbf{z}_l)$ via approximating the backward denoising process of $p(\mathbf{z}_h | \mathbf{z}_t, \mathbf{z}_l)$. Note that, for super-resolution, the denoising process is conditioned on the low-resolution input, \mathbf{z}_l , as well. This can

6 M. Noroozi et al.

be estimated by the function $\hat{\mathbf{z}}_{\theta}(\mathbf{z}_t, \mathbf{z}_l, \lambda_t)$ parametrized by a neural network. We can train this function via a weighted mean square error loss.

$$\underset{a}{\operatorname{argmin}} \ \mathbb{E}_{\epsilon,t}[\omega(\lambda_t) || \hat{\mathbf{z}}_{\theta}(\mathbf{z}_t, \mathbf{z}_l, \lambda_t) - \mathbf{z}_h ||_2^2]$$
(2)

over uniformly sampled times $t \in [0, 1]$ and $\mathbf{z}_t = \alpha_t \mathbf{z}_h + \sigma_t \epsilon$, $\epsilon \sim \mathcal{N}(0, I)$. There are several choices of weighting function $\omega(\lambda_t)$. We use the so called **v** parameterization [26], $(1 + \frac{\alpha_t^2}{\sigma_t^2})$, throughout this paper.

The inference process from a trained model involves a series of sequential calls, *i.e.* steps, of $\hat{\mathbf{z}}_{\theta}$, starting from $\mathbf{z}_1 \sim \mathcal{N}(0, I)$, where the quality of the generated image improves monotonically with the number of steps as shown in the qualitative examples of Fig .1 and quantitative results of Fig. 3. Several methods have been proposed to reduce the number of required steps at inference time [18, 19, 28]. Here, we use the widely used DDIM sampler [28], and yet see that the performance drops drastically with an extremely low number of steps. In the following, we introduce scale distillation to alleviate this shortcoming.

3.2 Scale distillation

The complexity of the image super-resolution task increases with the scale factor (SF). For example, a model trained for a lower SF $(e.g. \times 2)$ takes as input a less degraded image compared to a larger SF $(e.g. \times 4)$. Therefore, a diffusion model trained for $\times 2$ magnification should require fewer inference steps to solve the HR image generation task compared to a model trained for the x4 scale factor.

To alleviate the training complexity for larger scale factors, we build on this observation and propose a progressive scale distillation training strategy. In particular, we start by training a teacher for a lower SF that takes a less degraded image as input. We then use its prediction as a target to train the model for a higher factor as a student.

Let N be the target SF of interest. Standard training involves making pairs of low and high-resolution images, where the low-resolution image is smaller than the HR image by a factor of 1/N. The common approach for generating the training pairs is to gather a set of high-resolution images, perform synthetic degradation to obtain the corresponding low-resolution image and train a model that directly performs $\times N$ magnification [22, 32, 34] using eq. 2. Instead, we start by training a standard diffusion-based teacher for a lower SF, using a less degraded LR image, *e.g.* 2/N, as input and use its prediction to train the student.

More precisely, Let us assume $\hat{\mathbf{z}}_{\phi}$, $\hat{\mathbf{z}}_{\theta}$ be the teacher and student denoising models parameterized by ϕ , θ respectively. To train the student for a factor of N, we generate two degraded images for a given high-resolution image with factors 1/N, 2/N, with latent representations denoted by \mathbf{z}_l , \mathbf{z}'_l respectively. That means \mathbf{z}'_l is less degraded compared to \mathbf{z}_l . Similar to the standard diffusion model training, we sample random noise at t and add it to the high-resolution image to obtain \mathbf{z}_t . The scale distillation loss will be:



Fig. 2: Training pipeline of proposed scale distillation. For a given HR image (*e.g.* size 512×512) shown in green, we generate two degraded versions with factors of 2/N, 1/N (*e.g.* sizes 256×256 and 128×128), shown in yellow and red respectively. Both degraded images are resized back via bicubic upsampling to 512×512 to be used as input to the encoder, which projects them to $4 \times 64 \times 64$ tensors. The less and more degraded LR image is used as input to the teacher and student respectively via concatenation with the noisy version of the HR image, *i.e.* \mathbf{z}_t . The teacher's output is used as the target for training the student. Note that the teacher is first trained independently for a smaller magnification scale and then frozen during student training.

$$\underset{\theta}{\operatorname{argmin}} \quad \mathbb{E}_{\epsilon,t}[\omega(\lambda_t) || \hat{\mathbf{z}}_{\theta}(\mathbf{z}_t, \mathbf{z}_l, \lambda_t) - \hat{\mathbf{z}}_{\phi}(\mathbf{z}_t, \mathbf{z}'_l, \lambda_t) ||_2^2] \tag{3}$$

where the teacher is trained for N/2 magnification and frozen, and the student is initialized with the teacher's weights before the training. Note that we are using the latent diffusion framework that allows exactly the same architecture and input shapes for both the teacher and the student. Although the input lowresolution images for the student and teacher are of different sizes, they are both resized to a fixed size and fed to the encoder, which projects them to a tensor with a fixed size of $4 \times 64 \times 64$. Fig. 2 illustrates the proposed scale distillation process.

The idea of scale distillation is in line with that of progressive temporal distillation [26]. While a standard denoising model would only use the final image as the target irrespective of the sampled time step t (see Eq. 2), both scale and progressive temporal distillation rely on the teacher to provide a supervisory signal specific for step t (see Eq. 3). In this way, the supervisory signal is attuned to the specific denoising step, providing stable and consistent supervision at every denoising step. Fig. 3 provides empirical support for our hypothesis. We observe a significant gap between the distilled models from $\times 2$ to $\times 4$ and $\times 2$ to $\times 8$ compared to the models that are directly trained for $\times 4$ and $\times 8$, respectively.



Fig. 3: FID vs. number of DDIM steps on the DIV2K validation set obtained through bicubic degradation using SD for $\times 4$ and $\times 8$ magnifications trained with scale distillation and standard training. We use $\times 2 \rightarrow \times 4$ scale distillation for $\times 4$ and $\times 2 \rightarrow \times 4 \rightarrow \times 8$ for $\times 8$, and compare with the standard training directly for $\times 4$ and $\times 8$ respectively. All results are obtained using the original SD decoder. The model trained with scale distillation outperforms the standard training with large margin when using fewer steps for $\times 4$. The gap between scale distillation and the standard training is significantly higher for small $\times 8$ and remains steady for large numbers steps.

The gap is especially striking when evaluated with few inference steps and, as expected, shrinks as the number of steps increases and quality saturates.

Similar to the temporal progressive distillation [26], the proposed scale distillation process can be applied iteratively with higher scale factors at each training step. The first student is initialized from scratch and trained on the raw data, similar to the standard training. Consequently, this student becomes the new teacher for training the next scale factor. In this paper, we consider three distillation steps up to the scale factor of $\times 8$ starting from $\times 2$, *i.e.* $\times 2 \rightarrow \times 4 \rightarrow \times 8$. As it is shown in Fig. 3, scale distillation is significantly more effective for $\times 8$ magnification where the LR image is of even lower quality, thereby reinforcing the importance of our proposed progressive scale training strategy.

3.3 Decoder fine-tuning

While scale distillation improves the one-step inference noticeably, there is still a gap between the one-step model and the saturated performance with a larger number of steps, see Fig. 3. To fill this gap, we propose to fine-tune the decoder on top of the frozen one-step diffusion model resulting from scale distillation. That is, after training the diffusion model, we freeze the U-Net, apply one DDIM step for a given LR image, and use it as input to fine-tune the decoder for the SR task. We use the original loss that has been used for training the autoencoder [22]. Importantly, this fine-tuning strategy with the U-Net in place is only possible with a diffusion model that can work properly with one step as enabled by our scale distillation approach; see Table. 3. We empirically show that the combination of our scale distillation approach with decoder fine-tuning yields a one-step model that can readily compete with models requiring a large number of inference steps.

Implementation details. We use Stable diffusion v1.5 as our backbone and initialize our teacher with the text-to-image model. We use our own implementation of the v-parameterization with a cosine schedule. We use 4 A100 GPUs for all our experiments and train with a batch size of 60 with a gradient accumulation factor of 4.

4 Experiments

In this section, we evaluate our YONOS-SR against other methods targeting real image super-resolution at the standard $\times 4$ scale factor in Sec. 4.1 and demonstrate that our proposed scale distillation approach generalizes to higher scale factors of $\times 8$ in Sec. 4.2. We then provide qualitative results for $\times 4$ and $\times 8$ in Sec. 4.3. Finally, we perform ablation studies to highlight the role of our main contributions in Sec. 4.4.

4.1 Evaluation on real image super resolution

We begin by evaluating the performance of our proposed YONOS-SR model in the standard real image super-resolution setting targeting $\times 4$ scale factor.

Datasets. Following previous work (*e.g.* [2,32,34,41]), we use DIV2K [1], DIV8K [7], Flickr2k [30], OST [36] and a subset of 10K images from FFHQ training set [13] to train our model. We adopt the Real-ESRGAN [34] degradation pipeline to generate synthetic LR-HR pairs.

We then evaluate our model on both synthetic and real datasets. Similar to [32], we use 3K LR-HR (128 \rightarrow 512) pairs synthesized from the DIV2K validation set using the Real-ESRGAN degradation pipeline as our synthetic dataset. We also report results on the standard DIV2K validation split with bicubic degradations for completeness. For the real dataset, we use 128 × 128 center crops from the RealSR [11], DRealSR [38] and DPED-iphone [10] datasets.

Evaluation metrics. We evaluate using various perceptual and image quality metrics, including LPIPS [43], FID [9] (where applicable), as well as the no-reference image quality metric, MUSIQ [14]. For the synthetic datasets, we also report standard PSNR and SSIM metrics, for reference.

Baselines. As the main contribution of our paper targets improving the inference process of diffusion-based super-resolution, our main points of comparison are diffusion-based SR models, including the recent StableSR model [32], ReshShift [40], and the original LDM model [22]. For completeness, we also include comparison to other non-diffusion-based baselines, including; RealSR [11], BSRGAN [41], RealESRGAN [34], DASR [16] and FeMaSR [2].

10 M. Noroozi et al.

Datasets	Metrics	RealSR	BSRGAN	DASR	Real-ESRGAN+	FeMaSR	LDM	ResShift	StableSR	YONOS (ours)
DIV2K Valid RealESRGAN degradations	$FID \downarrow$	49.49	44.22	49.16	37.64	35.87	26.47	30.45	24.44	21.86
	LPIPS \downarrow	0.5276	0.3351	0.3543	0.3112	0.3199	0.2510	0.3076	0.3114	0.2310
	$PSNR \uparrow$	24.62	24.58	24.47	24.28	23.06	23.32	24.62	23.26	24.74
	$SSIM \uparrow$	0.5970	0.6269	0.6304	0.6372	0.5887	0.5762	0.6210	0.5726	0.6428
	MUSIQ ↑	28.57	61.19	55.19	61.05	60.83	62.27	63.58	65.92	70.30
DIV2K Valid	LPIPS \downarrow	-	0.2364	0.1696	0.2284	-	0.2323	0.1775	0.2580	0.1703
bicubic degradations	$PSNR \uparrow$	-	27.32	28.55	26.65	-	25.49	27.24	21.90	26.26
RealSR	LPIPS \downarrow	0.3570	0.2656	0.3134	0.2709	0.2937	0.3159	0.3279	0.3002	0.2479
	MUSIQ ↑	38.26	63.28	41.21	60.36	59.06	58.90	59.87	65.88	69.21
DRealSR	LPIPS \downarrow	0.3938	0.2858	0.3099	0.2818	0.3157	0.3379	0.3870	0.3284	0.2721
	MUSIQ ↑	26.93	57.16	42.41	54.26	53.71	53.72	54.13	58.51	66.26
DPED-iphone	MUSIQ ↑	45.60	45.89	32.68	42.42	49.95	44.23	38.59	50.48	59.45
-	# STEPS 1	-	-	-	-	-	200	4	200	1

Table 1: Comparison to baselines. Results in **Red** and Blue correspond to best and second best results, resp. Cells with – indicate that there were no previously reported results using the considered baseline and corresponding metric.

Results. Results summarized in Tab. 1 show that YONOS-SR outperforms all other diffusion-based SR methods, while using **only one inference step**, whereas other alternatives use 200 inference steps. These results highlight the efficiency of YONOS-SR in reducing the number of steps to one without compromising performance but indeed improving it further. Also, our model outperforms all considered baselines in 5 out of 7 metrics on the synthetic data and all comparison points on the real datasets.

4.2 Generalization to higher scale factors

We now evaluate the generalization capability of our proposed scale distillation approach. To this end, we train our YONOS-SR model with one more iteration of scale distillation, thereby going from a model capable of handling ×4 magnifications to ×8 magnifications. We then fine-tune the decoder on top of the one-step ×8 diffusion model. To evaluate this model, we follow recent work [3], and evaluate on the same subset of ImageNet and FFHQ for ×8 magnification, *i.e.* $64 \times 64 \rightarrow 512 \times 512$. In particular, we select the same 1k subset of ImageNet test set by first ordering the 10k images by name and then selecting the 1k subset via interleaved sampling, *i.e.* using images of index 0, 10, 20, etc. To obtain the LR-HR pairs, we use ×8 average pooling degradations. In the case of FFHQ, we use the first 1k images of the validation set. We also evaluate using the same metrics and baselines reported in this recent work [3].

The results summarized in Tab. 2 demonstrate that our proposed one-step method generalizes well to higher scale factors, where it is able to achieve good results in terms of FID and LPIPS scores, which are known to better align with human observation, especially at higher magnification factors [24]. Notably, unlike baselines, our model has not been trained on ImageNet data. We use only 10k images of FFHQ in our training set.

4.3 Qualitative evaluation

In addition to extensive quantitative evaluations, we qualitatively compare onestep YONOS-SR with 200-step StableSR and standard diffusion-based SR (SD-

YONOS-SR 11



Fig. 4: Qualitative comparison on the validation set of DIV2K dataset: (a) 200-step StableSR (b) 200-step standard SD-SR (c) 1-step YONOS(ours) (d) the ground truth. SD-SR represents the standard Stable Diffusion-based SR model. 200-step StableSR and SD-SR tend to over-sharpen, adding artifacts that do not match the ground truth content. Our SR images match the most with the corresponding ground truth image; see the faces, Pepsi, and crocodile textures in the first, second, and third rows, respectively. The images are best seen in a display and zoomed in.

SR) in Fig. 4. Our method generates the closest SR images to the ground truth in terms of detailed textures while taking only **1-step** during the inference. These observations are in line with the numerical superiority of our method in the quantitative evaluations above.

As it is clearly demonstrated in Fig. 3, scale distillation is even more effective for $\times 8$ compared to $\times 4$ magnification. As a qualitative support, we compare the model trained directly for $\times 8$ magnification without scale distillation to our model trained with three iterations of scale distillation $\times 2 \rightarrow \times 4 \rightarrow \times 8$ in Fig. 5. Again, we use the validation set of DIV2K dataset. In line with the numerical analyses in Fig. 3, we observe that the model trained with scale distillation out-



Fig. 5: Qualitative comparison on the validation set of DIV2K dataset for $\times 8$ magnification when the model is trained directly for $\times 8$ magnification without scale distillation (top row) and with three iterations of scale distillation $\times 2 \rightarrow \times 4 \rightarrow \times 8$ (bottom row). We show the input LR image results with 1, 4, and 64 steps using the original decoder and the corresponding HR image for both models. The model trained with scale distillation outperforms the standard training with high margins. Specifically, due to poor LR input, the standard training fails to recover the relevant content. The images are best seen in a display and zoomed in.

		Imagene	et	FFHQ				
	$\overline{\text{FID}}\downarrow$	LPIPS \downarrow	$\mathrm{PSNR}\uparrow$	$\overline{\mathrm{FID}}\downarrow$	LPIPS \downarrow	$PSNR\uparrow$		
LDPS	61.09	0.475	23.21	36.81	0.292	28.78		
GML-DPS [23]	60.36	0.456	23.21	41.65	0.318	28.50		
PSLD [23]	60.81	0.471	23.17	36.93	0.335	26.62		
LDIR [8]	63.46	0.480	22.23	36.04	0.345	25.79		
P2L [3]	51.81	0.386	23.38	31.23	0.290	28.55		
YONOS (ours)	34.59	0.241	22.80	21.41	0.161	26.08		

Table 2: Comparison to baselines on ImageNet subset with x8 magnification factor. The results for other methods are taken from [3].

performs the standard training in terms of recovering the corresponding content and details. Note that, the problem of $\times 8$ magnification is of significantly higher complexity compared to $\times 4$ due to poor LR input. Notable for these $\times 8$ qualitative evaluations we use the original decoder (*i.e.* these results are obtained before the decoder finetunig stage) to emphasize the impact of scale distillation.

4.4 Ablation study

We now study the impact of the various components introduced in our work. To this end, we use the standard DIV2K validation set with $\times 4$ low-resolution images obtained through bicubic degradation [1]. We use the FID metric as it is a standard metric for assessing the quality of generative models. Our initial investigation also revealed that FID correlates the most with the human evaluation of the generated images. The validation set of the DIV2K dataset includes only 100 samples. To obtain more reliable FID scores, we extract 30 random 128×128 patches and their corresponding 512×512 HR counterparts from each image in the standard DIV2K bicubic validation set, resulting in a total of 3k LR-HR pairs. For completeness, we also report LPIPS, PSNR, and SSIM scores.

Impact of scale distillation. We begin by evaluating the impact of our proposed scale distillation on speeding up inference time. To this end, we run two stable diffusion (SD) models trained for $\times 4$ super-resolution (SR), with various numbers of inference steps. The first model is a standard SD super-resolution model trained directly for target $\times 4$ super-resolution (*i.e.* SD-SR), while the second model is trained with our proposed scale distillation from $\times 2$ magnification to $\times 4$. We use the same model, training set, and degradation pipeline in training both models. The only difference is the use of our scale distillation in the later model. Specifically, we start with training a teacher for $\times 2$ magnification using raw data as a denoising target. We use the $\times 2$ model as a frozen teacher and use its prediction to train a student for $\times 4$ magnification. The results summarized in Fig. 3 speaks decisively in favor of our scale distillation performs significantly better than direct $\times 4$ training when using only one step.

Scale distillation outperforms the standard training more significantly for $\times 8$ magnification where we perform three training iterations for scale distillation, *i.e.* $\times 2 \rightarrow \times 4 \rightarrow \times 8$. One reason for the larger gap for $\times 8$ magnification is that the SR task is more ambiguous for $\times 8$ magnification due to lower quality input.

As a result, the model benefits more from the more simplified supervisory signal obtained from scale distillation. Note that we use the original SD decoder (i.e. no decoder finetuning) for this experiment to analyze the impact of the scale distillation independently of decoder fine-tuning.

Impact of decoder fine-tuning. One of the direct consequences of having a diffusion model that can yield good results in one denoising step is that it allows for decoder fine-tuning with the U-Net in place, as it will directly give a good starting point to the decoder. To validate the importance of the input given to the decoder prior to finetuning and, thereby, the importance of YONOS-SR, we experiment with the standard SD-SR model and our scale distillation model. In both cases, we freeze the U-Net and only allow the models to do 1 denoising step. We the

Decoder	Orig	ginal	Fine-tuned		
Scale distillation	X	\checkmark	X	\checkmark	
FID ↓	27.93	23.96	16.26	15.54	
\downarrow LPIPS \downarrow	0.227	0.207	0.163	0.159	
4 PSNR \uparrow	25.94	26.24	25.73	26.30	
SSIM \uparrow	0.711	0.714	0.713	0.727	
FID ↓	102.92	66.90	41.54	28.47	
LPIPS↓	0.541	0.403	0.305	0.243	
$^{\circ}$ PSNR \uparrow	21.08	24.46	21.53	23.96	
$SSIM \uparrow$	0.541	0.647	0.528	0.632	

Table 3: Role of scale distillation and decoder fine-tuning. All results reported here are obtained with 1 inference step.

models to do 1 denoising step. We then feed their output to the decoder and fine-tune it following the same loss used in the original stable diffusion model [22].

The results summarized in Tab. 3 validate the importance of having a good initial input from the diffusion model prior to decoder fine-tuning. The left chunk shows that the model trained with scale distillation outperforms the standard training with a good margin when using the original decoder, indicating that the scale distillation results in a U-Net that provides a higher quality input for the decoder. Moreover, as we can see in the right chunk of Tab. 3, fine-tuning the decoder on top of both 1-step models improves the performance. However, the model with scale distillation yields significantly better results than the standard SD-SR directly trained for the target magnification. Once again, the impact of scale distillation is more sensible for $\times 8$ magnification than $\times 4$, which highlights the importance of our approach in such difficult settings. Importantly, this fine-tuning strategy is not computationally feasible with diffusion models that require many denoising steps to give a reasonable starting point for the decoder.

5 Conclusion

In summary, in this paper, we introduced the first **fast** stable diffusion-based super-resolution method. To achieve this, we introduced scale distillation, an approach that allows us to tackle the SR problem in as little as one step. Having a fast diffusion model allowed us to directly fine-tune the decoder, which we show yields state-of-the-art results, even at high magnification factors and only using a single step. We hope that the proposed distillation approach could be adapted for other inverse imaging problems (*e.g.* image inpainting), which we believe is an interesting direction for future research.

References

- Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: IEEE Conference on Computer Vision and Pattern Recognition - Workshops (2017)
- Chen, C., Shi, X., Qin, Y., Li, X., Han, X., Yang, T., Guo, S.: Real-world blind super-resolution via feature matching with implicit high-resolution priors. In: ACM International Conference on Multimedia (2022)
- Chung, H., Ye, J.C., Milanfar, P., Delbracio, M.: Prompt-tuning latent diffusion models for inverse problems. In: arXiv preprint arXiv: 2310.01110 (2023)
- 4. Dong, C., Loy, C.C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: European Conference on Computer Vision (2014)
- Fritsche, M., Gu, S., Timofte, R.: Frequency separation for real-world superresolution. In: IEEE International Conference on Computer Vision - Workshops (2019)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances on Neural Information Processing Systems (2014)
- Gu, S., Lugmayr, A., Danelljan, M., Fritsche, M., Lamour, J., Timofte, R.: Div8k: Diverse 8k resolution image dataset. In: IEEE International Conference on Computer Vision - Workshops (2019)
- He, L., Yan, H., Luo, M., Luo, K., Wang, W., Du, W., Chen, H., Yang, H., Zhang, Y.: Iterative reconstruction based on latent diffusion model for sparse data reconstruction. In: arXiv preprint arXiv:2307.12070 (2023)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances on Neural Information Processing Systems (2017)
- Ignatov, A., Kobyshev, N., Timofte, R., Vanhoey, K., Gool, L.V.: Dslr-quality photos on mobile devices with deep convolutional networks. In: IEEE International Conference on Computer Vision (2017)
- 11. Ji, X., Cao, Y., Tai, Y., Wang, C., Li, J., Huang, F.: Real-world super-resolution via kernel estimation and noise injection. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (2020)
- Jolicoeur-Martineau, A., Li, K., Piché-Taillefer, R., Kachman, T., Mitliagkas, I.: Gotta go fast when generating data with score-based models. In: arXiv preprint arXiv:2105.14080 (2021)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
- 14. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yan, F.: Musiq: Multi-scale image quality transformer. In: IEEE International Conference on Computer Vision (2021)
- 15. Liang, J., Zhang, K., Gu, S., Van Gool, L., Timofte, R.: Flow-based kernel prior with application to blind superresolution. In: IEEE Conference on Computer Vision and Pattern Recognition (2021)
- Liang, J., Zeng, H., Zhang, L.: Efficient and degradation-adaptive network for real-world image super-resolution. In: European Conference on Computer Vision (2022)
- Liu, A., Liu, Y., Gu, J., Qiao, Y., Dong, C.: Blind image superresolution: A survey and beyond. In: arXiv preprint arXiv:2107.03055 (2021)

- 16 M. Noroozi et al.
- Lu, C., Zhou, Y., Bao, F., Chen, J., LI, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In: Advances on Neural Information Processing Systems (2022)
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. In: arxiv prepring arxiv: 2211.01095 (2023)
- 20. Maeda, S.: Unpaired image super-resolution using pseudo-supervision. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
- Meng, C., Rombach, R., Gao, R., Kingma, D., Ermon, S., Ho, J., Salimans, T.: On distillation of guided diffusion models. In: IEEE Conference on Computer Vision and Pattern Recognition (2023)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE Conference on Computer Vision and Pattern Recognition (2022)
- Rout, L., Raoof, N., Daras, G., Caramanis, C., and Sanjay Shakkottai, A.G.D.: Solving linear inverse problems provably via posterior sampling with latent diffusion models. In: NeurIPS (2023)
- Sahak, H., Watson, D., Saharia, C., Fleet, D.: Denoising diffusion probabilistic models for robust image super-resolution in the wild. In: arXiv preprint arXiv: 2302.07864 (2023)
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image superresolution via iterative refinement. preprint arXiv: 2104.07636 (2021)
- Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. In: International Conference on Learning Representations (2022)
- 27. Shocher, A., Cohen, N., Irani, M.: "zero-shot" superresolution using deep internal learning. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021)
- Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. arXiv preprint arXiv:2303.01469 (2023)
- Timofte, R., Agustsson, E., Gool, L.V., Yang, M., Zhang, L.: Ntire 2017 challenge on single image super-resolution: Methods and results. In: IEEE Conference on Computer Vision and Pattern Recognition - Workshops (2017)
- Wan, Z., Zhang, B., Chen, D., Zhang, P., Chen, D., Liao, J., Wen, F.: Bringing old photos back to life. In: IEEE Conference on Computer Vision and Pattern Recognition (2020)
- Wang, J., Yue, Z., Zhou, S., Chan, K.C., Loy, C.C.: Exploiting diffusion prior for real-world image super-resolution. In: arXiv preprint arXiv:2305.07015 (2023)
- Wang, L., Wang, Y., Dong, X., Xu, Q., Yang, J., An, W., Guo, Y.: Unsupervised degradation representation learning for blind superresolution. In: IEEE Conference on Computer Vision and Pattern Recognition (2021)
- Wang, X., Xie, L., Dong, C., Shan, Y.: Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data. In: IEEE International Conference on Computer Vision - Workshops (2021)
- Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image superresolution by deep spatial feature transform. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)
- Wang, X., Yu, K., Dong, C., Loy, C.C.: Recovering realistic texture in image superresolution by deep spatial feature transform. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)

- Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ES-RGAN: Enhanced super-resolution generative adversarial networks. In: European Conference on Computer Vision - Workshops (2018)
- Wei, P., Xie, Z., Lu, H., Zhan, Z., Ye, Q., Zuo, W., Lin, L.: Component divideand-conquer for real-world image super-resolution. In: European Conference on Computer Vision (2020)
- Yan, Y., Liu, C., Chen, C., Sun, X., Jin, L., Peng, X., Zhou, X.: Fine-grained attention and feature-sharing generative adversarial networks for single image superresolution. In: IEEE Transactions on Multimedia (2021)
- 40. Yue, Z., Wang, J., Change Loy, C.: Resshift: Efficient diffusion model for image super-resolution by residual shifting. In: NeurIPS (2023)
- Zhang, K., Liang, J., Van Gool, L., Timofte, R.: Designing a practical degradation model for deep blind image super-resolution. In: IEEE International Conference on Computer Vision (2021)
- 42. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: IEEE International Conference on Computer Vision (2023)
- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: IEEE Conference on Computer Vision and Pattern Recognition (2018)
- 44. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: IEEE International Conference on Computer Vision (2017)