

Explorable Super Resolution

Yuval Bahat and Tomer Michaeli

Technion - Israel Institute of Technology, Haifa, Israel

{yuval.bahat@campus,tomer.m@ee}.technion.ac.il



Figure 1: **Exploring HR explanations to an LR image.** Existing SR methods (e.g. ESRGAN [26]) output only one explanation to the input image. In contrast, our explorable SR framework allows producing infinite different perceptually satisfying HR images, that all identically match a given LR input, when down-sampled. Please zoom-in to view subtle details.

Abstract

Single image super resolution (SR) has seen major performance leaps in recent years. However, existing methods do not allow exploring the infinitely many plausible reconstructions that might have given rise to the observed low-resolution (LR) image. These different explanations to the LR image may dramatically vary in their textures and fine details, and may often encode completely different semantic information. In this paper, we introduce the task of explorable super resolution. We propose a framework comprising a graphical user interface with a neural network backend, allowing editing the SR output so as to explore the abundance of plausible HR explanations to the LR input. At the heart of our method is a novel module that can wrap any existing SR network, analytically guaranteeing that its SR outputs would precisely match the LR input, when down-sampled. Besides its importance in our setting, this module is guaranteed to decrease the reconstruction error of any SR network it wraps, and can be used to cope with blur kernels that are different from the one the network was trained for. We illustrate our approach in a variety of use cases, ranging from medical imaging and forensics, to graphics.

1. Introduction

Single image *super resolution* (SR) is the task of producing a high resolution (HR) image from a single low

resolution (LR) image. Recent decades have seen an increasingly growing research interest in this task, peaking with the recent surge of methods based on deep neural networks. These methods demonstrated significant performance boosts, some in terms of achieving low reconstruction errors [4, 10, 14, 24, 12, 29, 11] and some in terms of producing photo-realistic HR images [13, 26, 23], typically via the use of generative adversarial networks (GANs) [5]. However, common to all existing methods is that they do not allow exploring the abundance of plausible HR explanations to the input LR image, and typically produce only a *single* SR output. This is dissatisfying as although these HR explanations share the same low frequency content, manifested in their coarser image structures, they may significantly vary in their higher frequency content, such as textures and small details (see e.g., Fig. 1). Apart from affecting the image appearance, these fine details often encode crucial semantic information, like in the cases of text, faces and even textures (e.g., distinguishing a horse from a zebra). Existing SR methods ignore this abundance of valid solutions, and arbitrarily confine their output to a specific appearance with its particular semantic meaning.

In this paper, we initiate the study of *explorable super resolution*, and propose a framework for achieving it through user editing. Our method consists of a neural network utilized by a graphical user interface (GUI), which allows the user to interactively explore the space of perceptually pleasing HR images that could have given rise to



Figure 2: **An example user editing process.** Our GUI allows exploring the space of plausible SR reconstructions using a variety of tools. Here, local scribble editing is used to encourage the edited region to resemble the user’s graphical input. Then the entire shirt area (red) is edited by encouraging its patches to resemble those in the source (yellow) region. At any stage of the process, the output is *perfectly consistent* with the input (its down-sampled version identically matches the input).




Figure 3: **Visually examining the likelihood of patterns of interest.** Given an LR image of a car license plate, we explore the possible valid SR reconstructions by attempting to manipulate the central digit to appear like any of the digits 0 – 9, using our imprinting tool (see Sec. 4). Though the ground truth HR digit was 1, judging by the ESRGAN [26] result (or by our pre-edited reconstruction) would probably lead to misidentifying it as 0. In contrast, our results when imprinting digits 0,1 and 8 contain only minor artifacts, thus giving them similar likelihood.

a given LR image. An example editing process is shown in Fig. 2. Our approach is applicable in numerous scenarios. For example, it enables manipulating the image so as to fit any prior knowledge the user may have on the captured scene, like changing the type of flora to match the capturing time and location, adjusting shades according to the capturing time of day, or manipulating an animal’s appearance according to whether the image was taken in the zebras or horses habitat. It can also help determine whether a certain pattern or object could have been present in the scene. This feature is invaluable in many settings, including in the forensic and medical contexts, exemplified in Figs. 3 and 4, respectively. Finally, it may be used to correct unpleasing SR outputs, which are common even with high capacity neural network models.

Our framework (depicted in Fig. 5) consists of three key ingredients, which fundamentally differ from the common




Figure 4: **Visually examining the likelihood of a medical pathology.** Given an LR shoulder X-ray image, we evaluate the likelihood of a Supraspinatus tendon tear, typically characterized by a less than 7mm Acromioclavicular distance (measured between the Humerus bone, marked red, and the Acromion bone above it). To this end, we attempt to imprint down-shifted versions (see Sec. 4) of the Acromion bone. Using the image quality as a proxy for its plausibility, we can infer a low chance of pathology, due to the artifacts emerging when forcing the pathological form (right image).

practice in SR. (i) We present a novel *consistency enforcing module* (CEM) that can wrap any SR network, analytically guaranteeing that its outputs identically match the input, when down-sampled. Besides its crucial role in our setting, we illustrate the advantages of incorporating this module into any SR method. (ii) We use a neural network with a control input signal, which allows generating diverse HR explanations to the LR image. To achieve this, we rely solely on an adversarial loss to promote perceptual plausibility, without using *any reconstruction loss* (e.g. L_1 or VGG) for promoting proximity between the network’s outputs and the ground-truth HR images. (iii) We facilitate the exploration process by creating a GUI comprising a *large set of tools*. These work by manipulating the network’s control signal so as to achieve various desired effects¹. We elaborate on those three ingredients in Secs. 2,3 and 4.

¹Our code and GUI are available online.

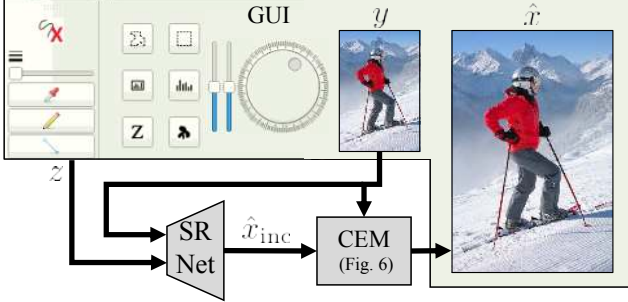


Figure 5: **Our explorable SR framework.** Our GUI allows interactively exploring the manifold of possible HR explanations for the LR image y , by manipulating the control signal z of the SR network. Our CEM is utilized to turn the network’s output \hat{x}_{inc} into a consistent reconstruction \hat{x} , presented to the user. See Secs. 2,3 and 4 for details.

1.1. Related Work

GAN based image editing Many works employed GANs for image editing tasks. For example, Zhu *et al.* [30] performed editing by searching for an image that satisfies a user input scribble, while constraining the output image to lie on the natural image manifold, learned by a GAN. Perarnau *et al.* [21] suggested to perform editing in a learned latent space, by combining an encoder with a conditional GAN. Xian *et al.* [28] facilitated texture editing, by allowing users to place a desired texture patch. Rott Shaham *et al.* [23] trained their GAN solely on the image to be edited, thus encouraging their edited output to retain the original image characteristics. While our method also allows traversing the natural image manifold, it is different from previous approaches in that it enforces the hard consistency constraint (restricting all outputs to identically match the LR input when down-sampled).

GAN based super-resolution A large body of work has shown the advantage of using *conditional GANs* (cGANs) for generating photo-realistic SR reconstructions [13, 26, 22, 25, 27]. Unlike classical GANs, cGANs feed the generator with additional data (*e.g.* an LR image) together with the stochastic noise input. The generator then learns to synthesize new data (*e.g.* the corresponding SR image) conditioned on this input. In practice, though, cGAN based SR methods typically feed their generator only with the LR image, without a stochastic noise input. Consequently, they do not produce diverse SR outputs. While we also use a cGAN, we do add a control input signal to our generator’s LR input, which allows editing its output to yield diverse results. Several cGAN methods for image translation did target outputs’ diversity by keeping the additional stochastic input [31, 3, 9], while utilizing various mechanisms for binding the output to this additional input. In our method,

we encourage diversity by simply removing the reconstruction losses that are used by all existing SR methods. This is made possible by our consistency enforcing module.

2. The Consistency Enforcing Module

We would like the outputs of our explorable SR method to be both perceptually plausible and consistent with the LR input. To encourage perceptual plausibility, we adopt the common practice of utilizing an adversarial loss, which penalizes for deviations from the statistics of natural images. To guarantee consistency, we introduce the *consistency enforcing module* (CEM), an architecture that can wrap any given SR network, making it inherently satisfy the consistency constraint. This is in contrast to existing SR networks, which do not perfectly satisfy this constraint, as they encourage consistency only *indirectly* through a reconstruction loss on the SR image. The CEM does not contain any learnable parameters and has many notable advantages over existing SR network architectures, on which we elaborate later in this section. We next derive our module.

Assume we are given a low resolution image y that is related to an unknown high-resolution image x through

$$y = (h * x) \downarrow_{\alpha}. \quad (1)$$

Here, h is a blur kernel associated with the point-spread function of the camera, $*$ denotes convolution, and \downarrow_{α} signifies sub-sampling by a factor α . With slight abuse of notation, (1) can be written in matrix form as

$$y = Hx, \quad (2)$$

where x and y now denote the vectorized versions of the HR and LR images, respectively, and the matrix H corresponds to convolution with h and sub-sampling by α . This system of equations is obviously under-determined, rendering it impossible to uniquely recover x from y without additional knowledge. We refer to any HR image \hat{x} satisfying this constraint, as *consistent* with the LR image y . We want to construct a module that can project any inconsistent reconstruction \hat{x}_{inc} (*e.g.* the output of a pre-trained SR network) onto the affine subspace defined by (2). Its consistent output \hat{x} is thus the minimizer of

$$\min_{\hat{x}} \|\hat{x} - \hat{x}_{\text{inc}}\|^2 \quad \text{s.t.} \quad H\hat{x} = y. \quad (3)$$

Intuitively speaking, such a module would guarantee that the low frequency content of \hat{x} matches that of the ground-truth image x (manifested in y), so that the SR network should only take care of plausibly reconstructing the high frequency content (*e.g.* sharp edges and fine textures).

Problems like (3) frequently arise in sampling theory (see *e.g.* [17]), and can be conveniently solved using a geometric viewpoint. Specifically, let us utilize the fact that




Figure 6: **CEM architecture.** The CEM, given by Eq. (8), can wrap any given SR network. It projects its output \hat{x}_{inc} onto the space of images that identically match input y when downsampled, thus producing a super-resolved image \hat{x} guaranteed to be consistent with y . See Sec. 2 for details.

$P_{\mathcal{N}(H)^\perp} = H^T(HH^T)^{-1}H$ is known to be the orthogonal projection matrix onto $\mathcal{N}(H)^\perp$, the subspace that is perpendicular to the nullspace of H . Now, multiplying both sides of the constraint in (3) by $H^T(HH^T)^{-1}$, yields

$$P_{\mathcal{N}(H)^\perp} \hat{x} = H^T(HH^T)^{-1}y. \quad (4)$$

This shows that we should strictly set the component of \hat{x} in $\mathcal{N}(H)^\perp$ to equal the right hand side of (4).

We are therefore restricted to minimize the objective by manipulating only the complementary component, $P_{\mathcal{N}(H)} \hat{x}$, that lies in the nullspace of H . Decomposing the objective into the two subspaces,

$$\|P_{\mathcal{N}(H)}(\hat{x} - \hat{x}_{\text{inc}})\|^2 + \|P_{\mathcal{N}(H)^\perp}(\hat{x} - \hat{x}_{\text{inc}})\|^2, \quad (5)$$

we see that $P_{\mathcal{N}(H)} \hat{x}$ only appears in the first term, which it minimizes when set to

$$P_{\mathcal{N}(H)} \hat{x} = P_{\mathcal{N}(H)} \hat{x}_{\text{inc}}. \quad (6)$$

Combining the two components from (4) and (6), and using the fact that $P_{\mathcal{N}(H)} = I - H^T(HH^T)^{-1}H$, we get that

$$\begin{aligned} \hat{x} &= P_{\mathcal{N}(H)} \hat{x} + P_{\mathcal{N}(H)^\perp} \hat{x} \\ &= (I - H^T(HH^T)^{-1}H) \hat{x}_{\text{inc}} + H^T(HH^T)^{-1}y. \end{aligned} \quad (7)$$

To transform (7) into a practical module that can wrap any SR architecture with output \hat{x}_{inc} , we need to replace the impractical multiplication operations involving the very large H matrix, with their equivalent operations: convolutions, downsampling and upsampling. To this end, we observe that since H corresponds to convolution with h followed by sub-sampling, H^T corresponds to up-sampling followed by convolution with a mirrored version of h , which we denote by \tilde{h} . The multiplication by $(HH^T)^{-1}$ can then be replaced by convolving with a filter k , constructed by computing the inverse of the filter $(h * \tilde{h}) \downarrow_\alpha$ in the Fourier domain. We thus have that

$$\hat{x} = \hat{x}_{\text{inc}} - \tilde{h} * [k * (h * \hat{x}_{\text{inc}}) \downarrow_\alpha] \uparrow_\alpha + \tilde{h} * (k * y) \uparrow_\alpha. \quad (8)$$

	2×	3×	4×
EDSR [14]	35.97	32.27	30.30
EDSR+CEM	36.11	32.36	30.37

Table 1: **Wrapping a pre-trained SR network with CEM.** Mean PSNR values over the BSD100 dataset [15], super-resolved by factors 2, 3 and 4. Merely wrapping the network with our CEM can only improve the reconstruction error, as manifested by the slight PSNR increase in the 2nd row.

Thus, given the blur kernel² h , we can calculate the filters appearing in (8) and hardcode their non-learned weights into our CEM, which can wrap any SR network, as shown in Fig. 6 (see Supplementary for padding details). Before proceeding to incorporate it in our scheme, we note the CEM is beneficial for *any* SR method, in the following two aspects.

Reduced reconstruction error Merely wrapping a pre-trained SR network with output \hat{x}_{inc} by our CEM, can only decrease its reconstruction error w.r.t. the ground-truth x , as

$$\begin{aligned} \|\hat{x}_{\text{inc}} - x\|^2 &\geq \|P_{\mathcal{N}(H)}(\hat{x}_{\text{inc}} - x)\|^2 \stackrel{(a)}{=} \|P_{\mathcal{N}(H)}(\hat{x} - x)\|^2 \\ &\stackrel{(b)}{=} \|\hat{x} - x\|^2. \end{aligned} \quad (9)$$

Here, (a) is due to (6), and (b) follows from (4), which implies that $\|P_{\mathcal{N}(H)^\perp}(\hat{x} - x)\|^2 = 0$ (as $P_{\mathcal{N}(H)^\perp} \hat{x} = H^T(HH^T)^{-1}Hx = P_{\mathcal{N}(H)}x$). This is demonstrated in Tab. (1), which reports PSNR values attained by the EDSR network [14] with and without our CEM, for several different SR scale factors.

Adopting to a non-default down-sampling kernel Deep learning based SR methods are usually trained on LR images obtained by down-sampling HR images using a specific down-sampling kernel (usually the bicubic kernel). This constitutes a major drawback, as their performance significantly degrades when processing LR images corresponding to different kernels, as is the case with most realistic images [18, 24]. This problem can be alleviated using our CEM, that takes the down-sampling kernel as a parameter upon its assembly at test time. Thus, it can be used for adopting a given network that was pre-trained using some default kernel, to the kernel corresponding to the LR image. Figure 7 demonstrates the advantage of this approach on a real LR image, using the kernel estimation method of [2].

3. Editable SR Network

To achieve our principal goal of creating an explorable super resolution network, we develop a new framework, that comprises a GUI with a neural network backend (see Fig. 5). Our SR network differs from existing SR methods in two core aspects. First, it is capable of producing a

²We default to the bicubic blur kernel when h is not given.



Figure 7: **Accounting for a different blur kernel at test time.** The CEM can be used to adopt any SR network to a blur kernel other than the one it was trained on. We demonstrate this on a real LR image. Here, ESRGAN [26], which was trained with the bicubic kernel, produces blurry results. However, wrapping it with our CEM, with the kernel estimated by [2], results in a much sharper reconstruction. Corresponding kernels are visualized on each image.

wide variety of plausible and inherently consistent HR reconstructions \hat{x} , for any input LR image y . This is achieved thanks to the CEM, which omits the need to use any reconstruction loss (e.g. L_1 or VGG) for driving the outputs close on average to the corresponding ground truth training images. Such losses bias the outputs towards the average of all possible explanations to the LR image, and are thus not optimal for the purpose of exploration. Second, our network incorporates a control input signal z , that allows traversing this manifold of plausible images so as to achieve various manipulation effects on the output \hat{x} . Using our GUI, the user can either tune z manually to achieve simple effects (e.g. control the orientation and magnitude of gradients as in Fig. 8) or use any of our automatic manipulation tools that take as input e.g. user scribbles or some pattern to imprint (see Figs. 2-4, 10, 12).

We use the same network architecture as ESRGAN [26], but train it wrapped by the CEM, and by minimizing a loss function comprising four terms,

$$\mathcal{L}_{\text{Adv}} + \lambda_{\text{Range}} \mathcal{L}_{\text{Range}} + \lambda_{\text{Struct}} \mathcal{L}_{\text{Struct}} + \lambda_{\text{Map}} \mathcal{L}_{\text{Map}}. \quad (10)$$

Here, \mathcal{L}_{Adv} is an adversarial loss, which encourages the network outputs to follow the statistics of natural images. We specifically use a Wasserstein GAN loss with gradient penalty [6], and avoid using the relativistic discriminator [8] employed in ESRGAN, as it induces a sort of full-reference supervision. The second loss term, $\mathcal{L}_{\text{Range}}$, penalizes for pixel values that exceed the valid range $[0, 1]$, and thus helps prevent model divergence. We use $\mathcal{L}_{\text{Range}} = \frac{1}{N} \|\hat{x} - \text{clip}_{[0,1]} \{\hat{x}\}\|_1$, where N is the number of pixels. The

	Diversity (σ)	Percept. quality (NIQE)	Reconst. error (RMSE)
ESRGAN	0	3.5 ± 0.9	17.3 ± 7.2
ESRGAN with z	3.6 ± 1.7	3.7 ± 0.8	17.5 ± 6.9
Ours	7.2 ± 3.4	3.7 ± 0.9	18.2 ± 7.4

Table 2: **Quality and diversity of SR results.** We report diversity (standard deviation, higher is better), perceptual quality (NIQE [19], lower is better) and RMSE (lower is better), for $4\times$ SR on the BSD100 test set [15]. Values are measured over 50 different SR outputs per input image, produced by injecting 50 random, spatially uniform z inputs. Note that our model, trained without any full-reference loss terms, shows a significant advantage in terms of diversity, while exhibiting similar perceptual quality. See Supplementary for more details about this experiment.

last two loss terms, $\mathcal{L}_{\text{Struct}}$ and \mathcal{L}_{Map} , are associated with the control signal z . We next elaborate on the control mechanism and these two penalties.

3.1. Incorporating a Control Signal

As mentioned above, to enable editing the output image \hat{x} , we introduce a control signal z , which we feed to the network in addition to the input image y . We define the control signal as $z \in \mathbb{R}^{w \times h \times c}$, where $w \times h$ are the dimensions of the output image \hat{x} and $c = 3$, to allow intricate editing abilities (see below). To prevent the network from ignoring this additional signal, as reported for similar cases in [7, 16], we follow the practice in [20] and concatenate z to the input of each layer of the network, where layers with smaller spatial dimensions are concatenated with a spatially downsampled version of z . At test time, we use this signal to traverse the space of plausible HR reconstructions. Therefore, at train time, we would like to encourage the network to associate different z inputs to different HR explanations. To achieve this, we inject random z signals during training.

Incorporating this input signal into the original ESRGAN method already affects outputs diversity. This can be seen in Tab. 2, which compares the vanilla ESRGAN method (1st row) with its variant, augmented with z as described above (2nd row), both trained for additional 6000 generator steps using the original ESRGAN loss. However, we can obtain an even larger diversity. Specifically, recall that as opposed to the original ESRGAN, in our loss we use no reconstruction (full-reference) penalty that resists diversity. The effect can be seen in the 3rd row of Tab. 2, which corresponds to our model trained for the same number of steps³ using only the \mathcal{L}_{Adv} and $\mathcal{L}_{\text{Range}}$ loss terms. Note that

³Weights corresponding to z in the 2nd and 3rd rows' models are initialized to 0, while all other weights are initialized with the pre-trained




Figure 8: **Manipulating image gradients.** A user can explore different perceptually plausible textures (e.g. the Mandrill’s cheek fur), by manually adjusting the directions and magnitudes of image gradients using the control signal z .

all three models are on par in terms of perceptual quality.

Now that the output of our network strongly depends on z , we move on to make this dependence easy to manipulate by a user. To allow manual control, we want simple modifications to z to lead to variations that are intuitive to grasp. To allow optimization-based manipulations, we want to ensure that any plausible output image could be generated by some z . These two requirements are encouraged by the last two loss terms in (10), as we describe next.

Controlling image gradients with the $\mathcal{L}_{\text{Struct}}$ loss To allow intuitive manual control, we encourage spatially uniform perturbations to z to affect the spatial derivatives of the output image \hat{x} . Figure 8 demonstrates how manipulating the gradients’ magnitudes and directions, can be used to edit textures, like fur. We choose to associate the 3 channels of z with the 3 degrees of freedom of the structure tensor of \hat{x} , which is the 2×2 matrix defined by $\mathcal{S}_{\hat{x}} = \iint (\nabla \hat{x}(\xi, \eta)) (\nabla \hat{x}(\xi, \eta))^T d\xi d\eta$. We encode this link into our network through the loss term $\mathcal{L}_{\text{Struct}}$, which penalizes for the difference between desired structure tensors \mathcal{S}_d , determined by a randomly sampled, spatially uniform z , and the tensors corresponding to the actual network’s outputs, $\mathcal{S}_{\hat{x}}$ (see Supplementary for more details).

Facilitating optimization based editing via the \mathcal{L}_{Map} loss Denoting our network output by $\hat{x} = \psi(y, z)$, we would like to guarantee that ψ can generate every plausible image \hat{x} with some choice of z . To this end, we introduce the loss term $\mathcal{L}_{\text{Map}} = \min_z \|\psi(y, z) - x\|_1$, which penalizes for differences between the real natural image x , and its best possible approximation using some signal z . Within each training step, we first solve the internal minimization of \mathcal{L}_{Map} over z for 10 iterations, and then freeze this z for the minimization of all loss terms in (10). Note that in contrast to $\mathcal{L}_{\text{Struct}}$, which only involves spatially uniform z inputs, the minimization over z in \mathcal{L}_{Map} exposes our network to the entire z space during training, encouraging its mapping onto the entire natural image manifold. Figure 9 illustrates how incorporating the \mathcal{L}_{Map} loss term improves our model’s ability to generate plausible patterns.




Figure 9: **Effect of \mathcal{L}_{Map} .** Incorporating the \mathcal{L}_{Map} loss term into our model’s training introduces spatially non-uniform z inputs, and improves the model’s ability to cover the natural image manifold. We demonstrate this by comparing a model trained with vs. without this term, in the task of imprinting the digit ‘8’ on the license plate of Fig. 3.

3.2. Training details

We train our network using the 800 DIV2K training images [1]. We use the standard bicubic downsampling kernel, both for producing LR training inputs and as h in our CEM (Eq. (8)), and feed our network with 64×64 input patches, randomly cropped from these LR images. We initialize our network weights with the pre-trained ESRGAN [26] model, except for weights corresponding to input signal z , which we initialize to zero. We minimize the loss in (10) with $\lambda_{\text{Range}} = 5000$, $\lambda_{\text{Struct}} = 1$ and $\lambda_{\text{Map}} = 100$. We set the Wasserstein GAN gradient penalty weight to $\lambda_{\text{gp}} = 10$. We establish critic’s credibility before performing each generator step, by verifying that the critic correctly distinguishes fake images from real ones for 10 consecutive batches. We use a batch size of 48 and train for $\sim 80K$ steps.

4. Editing Tools

We incorporate our trained SR network described in Sec. 3 as the backend engine of an editing GUI, that allows manipulating the output SR image \hat{x} by properly shaping the control signal z . Our GUI comprises several tools that can be applied on selected regions of the image, which we achieve by manipulating only the corresponding regions of z (recall that the dimensions of z and \hat{x} are the same).

The most basic tool constitutes three knobs (top middle in Fig. 5) that manually control the triple channeled z signal in a spatially uniform manner, so as to affect the image structure tensor $\mathcal{S}_{\hat{x}}$ (Sec. 3.1). We make this tool user intuitive by binding the knobs with the eigenvalue decomposition of $\mathcal{S}_{\hat{x}}$, thus affecting image gradients. The user

ESRGAN model from the 1st row.

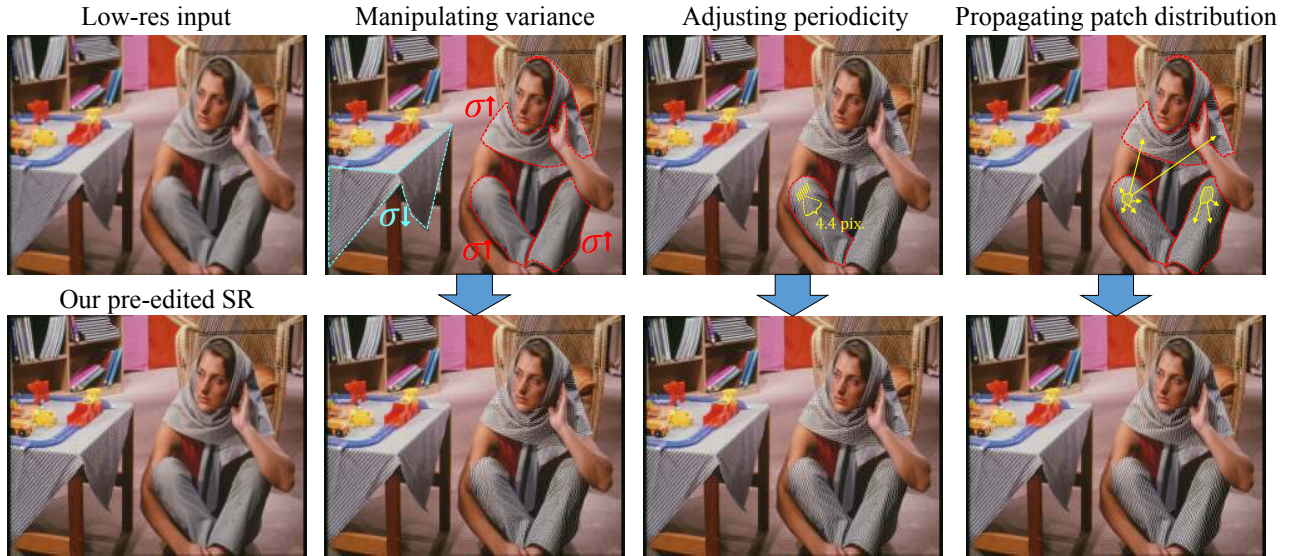


Figure 10: **Editing using tools that optimize over z .** Here we use our variance manipulation, periodicity, and patch dictionary tools. Each optimizes over z space to achieve a different objective. See Sec. 4 for more details.

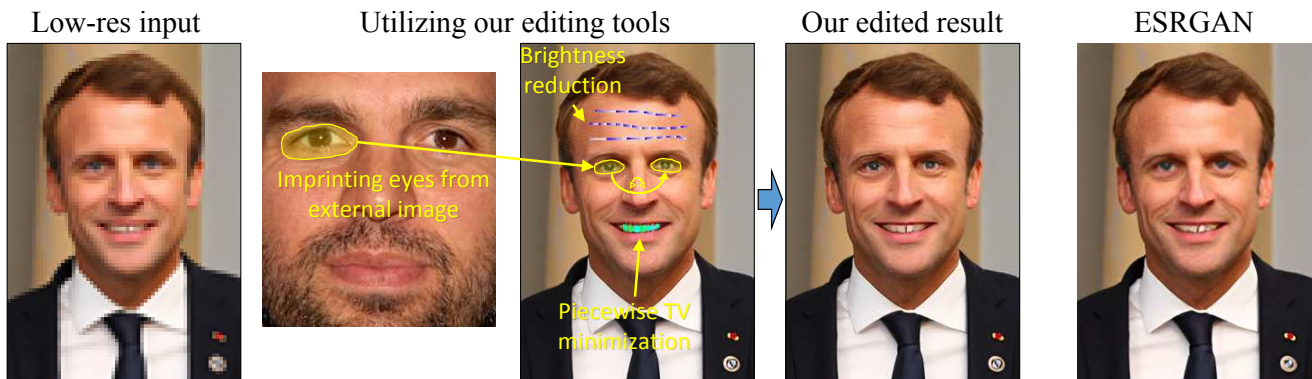


Figure 11: **Utilizing scribble editing tools.** Our tools allow regaining small details and producing crisper and more realistic images (compared to, *e.g.* ESRGAN [26]). Here we use the piecewise smoothness tool (TV minimization) to enhance edges between teeth, and restore forehead wrinkles using our brightness reduction pen. We imprint the eye region from an online available image of (brown-eyed) Mark Ruffalo to restore the crisp appearance of the blue eyes of Emmanuel Macron.

can control the orientation θ and magnitude λ_1 of prominent gradients, and the magnitude λ_2 in the perpendicular direction (see Fig. 8).

To allow a more diverse set of intricate editing operations, we also propose tools that optimize specific editing objectives (*e.g.* increasing local image variance). These tools invoke a gradient descent optimization process over z , whose goal is to minimize the chosen objective. This is analogous to traversing the manifold of perceptually plausible valid SR images (captured by our network), while remaining consistent with the LR image (thanks to the CEM).

Our GUI recomputes \hat{x} after each edit in nearly real-time, ranging from a couple of milliseconds for the basic

tools to 2-3 seconds for editing a 100×100 region using the z optimization tools (with an NVIDIA GeForce 2080 GPU). We next briefly survey and demonstrate these tools, and leave the detailed description of each objective to the supplementary.

Variance manipulation This set of tools searches for a z that decreases or increases the variance of pixel values. This is demonstrated in Fig. 10 for yielding a smoother table map and more textured trousers. A variant of this tool allows *locally* magnifying or attenuating exiting patterns.

User scribble Our GUI incorporates a large set of tools for imposing a graphical input by the user on the SR image.

A user first scribbles on the current SR image and chooses the desired color and line width. Our GUI then optimizes over z , searching for the image \hat{x} that satisfies the imposed input while lying on the learned manifold of valid perceptually plausible SR images. Figure 2 (1st editing operation) demonstrates how we use line scribbles for creating the desired shirt pattern. Variants of this tool (demonstrated in Fig. 11) enable increasing or decreasing brightness, as well as enforcing smoothness across a specific scribble mark, by minimizing local TV.

Imprinting This tool enables a user to impose content taken from an external image, or from a different region within the same image. After selecting the desired region, our GUI utilizes the CEM to combine the low frequency content corresponding to y with the high frequency content of the region to be imprinted. This results in a consistent image region, but which may not necessarily be perceptually plausible. We then invoke an optimization over z , which attempts to drive \hat{x} close to this consistent region, while remaining on the natural image manifold. Example usages are shown in Figs. 1 (right-most image), 3, 11 and 4, where the tool is used for subtle shifting of image regions.

Desired dictionary of patches This tool manipulates the edited region to comprise patches stemming from a desired patch collection. A user first marks a source region containing desired patches. This region may be taken either from the edited image or from an external image. Our GUI then optimizes z to encourage the target edited region to comprise the same desired patches, where we use patches of size 6×6 . Figures 2 (2nd editing operation) and 10 (right) show how we use this tool to propagate a desired local cloth pattern (yellow region) to the entire garment. Variants of this tool allow matching patches while ignoring their mean value or variance, to allow disregarding color and pattern magnitude differences between source and target patches.

Encouraging periodicity This tool enhances the periodic nature of the image along one or two directions marked by the user. The desired period length can be manually adjusted, as exemplified in Fig. 10 for the purpose of producing different stripe widths. Alternatively, a user can choose to enhance the most prominent existing period length, automatically estimated by our GUI.

Random diverse alternatives This tool offers an easy way to explore the image manifold. When invoked, it optimizes z to simultaneously produce N different SR image alternatives, by maximizing the L_1 distance between them, in pixel space. The user can then use each of the alternatives (or sub-regions thereof) as a baseline for further editing. A variant of this tool constrains the different alternatives to remain close to the current SR image.

The wide applicability of our method and editing tools is further demonstrated in Figs. 12 and 13 in two use cases.

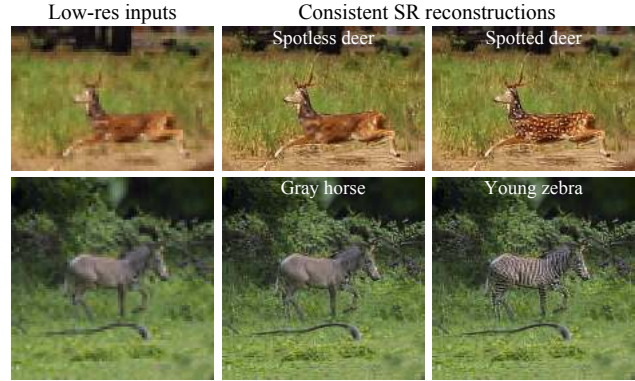


Figure 12: **More exploration examples.** Using our framework to explore possible SR images corresponding to different semantic contents, all consistent with the LR input.




Figure 13: **Correcting SR outputs.** Our method can be used to enhance SR results, relying on the user’s prior knowledge (e.g. the appearance of buildings and Alphabet characters). Please zoom-in to view subtle differences, and refer to the Supplementary for the editing processes.

Figure 12 demonstrates our method’s exploration capabilities, illustrating how utterly different semantic content can be generated for the same LR input. Figure 13 shows how our method’s editing capabilities can enhance SR outputs, by relying on the user’s external knowledge. Please see more editing process examples and results in Supplementary.

5. Conclusion

A low resolution image may correspond to many different HR images. However, existing SR methods generate only a single, arbitrarily chosen, reconstruction. In this work, we introduced a framework for explorable super resolution, which allows traversing the set of perceptually plausible HR images that are consistent with a given LR image. We illustrated how our approach is beneficial in many domains, including in medicine, surveillance, and graphics.

Acknowledgment This research was supported by the Israel Science Foundation (grant 852/17) and by the Technion Ollendorff Minerva Center.

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 126–135, 2017. 6
- [2] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. *arXiv preprint arXiv:1909.06581*, 2019. 4, 5
- [3] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pages 2172–2180, 2016. 3
- [4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *European conference on computer vision*, pages 184–199. Springer, 2014. 1
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1
- [6] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017. 5
- [7] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arxiv*, 2016. 5
- [8] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard gan. *arXiv preprint arXiv:1807.00734*, 2018. 5
- [9] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 3
- [10] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016. 1
- [11] Idan Kligvasser, Tamar Rott Shaham, and Tomer Michaeli. xunit: Learning a spatial activation function for efficient image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2433–2442, 2018. 1
- [12] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Fast and accurate image super-resolution with deep laplacian pyramid networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 1
- [13] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 1, 3
- [14] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 1, 4
- [15] David Martin, Charless Fowlkes, Doron Tal, Jitendra Malik, et al. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings of the IEEE International Conference on Computer Vision*, 2001. 4, 5
- [16] Michael Mathieu, Camille Couprie, and Yann LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015. 5
- [17] Tomer Michaeli and Yonina C. Eldar. *Optimization techniques in modern sampling theory*. Cambridge, UK: Cambridge Univ. Press, 2010. 3
- [18] Tomer Michaeli and Michal Irani. Nonparametric blind super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 945–952, 2013. 4
- [19] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a completely blind image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012. 5
- [20] Pablo Navarrete Michelini, Dan Zhu, and Hanwen Liu. Multi-scale recursive and perception-distortion controllable image super-resolution. In *Proceedings of the European Conference on Computer Vision Workshops*, 2018. 5
- [21] Guim Perarnau, Joost Van De Weijer, Bogdan Raducanu, and Jose M Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016. 3
- [22] Mehdi SM Sajjadi, Bernhard Scholkopf, and Michael Hirsch. Enhancenet: Single image super-resolution through automated texture synthesis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4491–4500, 2017. 3
- [23] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. *arXiv preprint arXiv:1905.01164*, 2019. 1, 3
- [24] Assaf Shocher, Nadav Cohen, and Michal Irani. zero-shot super-resolution using deep internal learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3118–3126, 2018. 1, 4
- [25] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 606–615, 2018. 3
- [26] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018. 1, 2, 3, 5, 6, 7
- [27] Yifan Wang, Federico Perazzi, Brian McWilliams, Alexander Sorkine-Hornung, Olga Sorkine-Hornung, and Christopher Schroers. A fully progressive approach to single-image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 864–873, 2018. 3

- [28] Wenqi Xian, Patsorn Sangkloy, Varun Agrawal, Amit Raj, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. Texturegan: Controlling deep image synthesis with texture patches. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 3
- [29] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [30] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision*, pages 597–613. Springer, 2016. 3
- [31] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 465–476, 2017. 3