

Single Image Super-Resolution via CNN Architectures and TV-TV Minimization

Marija Vella
mv37@hw.ac.uk
João F. C. Mota
j.mota@hw.ac.uk

Institute of Sensors, Signals and
Systems
Heriot-Watt University
Edinburgh, UK

Abstract

Super-resolution (SR) is a technique that allows increasing the resolution of a given image. Having applications in many areas, from medical imaging to consumer electronics, several SR methods have been proposed. Currently, the best performing methods are based on convolutional neural networks (CNNs) and require extensive datasets for training. However, at test time, they fail to impose consistency between the super-resolved image and the given low-resolution image, a property that classic reconstruction-based algorithms naturally enforce in spite of having poorer performance. Motivated by this observation, we propose a new framework that joins both approaches and produces images with superior quality than any of the prior methods. Although our framework requires additional computation, our experiments on Set5, Set14, and BSD100 show that it systematically produces images with better peak signal to noise ratio (PSNR) and structural similarity (SSIM) than the current state-of-the-art CNN architectures for SR.

Introduction

Image sensing technology is fundamentally limited by the physics of the acquisition process and often produces images with resolution below expectation. Super-resolution (SR) techniques overcome this by taking as input a low-resolution (LR) image and by outputting a high-resolution (HR) version of it. This, however, requires filling in the missing entries of the image, for which there are many possibilities. Predicting these entries thus requires making assumptions about the images, and different techniques make different assumptions. Landmark techniques include polynomial interpolation methods, which assume that images are spatially smooth and well approximated by bilinear or bicubic functions [6, 13], reconstruction methods, which assume that images have sparse representations in certain domains that are fixed [2, 32, 37] or that can be learned from training data [40, 41], and convolutional neural network (CNN) architectures [11, 19, 21], which assume that image patches share high-level features that can be learned by a CNN.

Indeed, the best performing SR methods are based on CNNs [11, 19, 21]. Although they produce outstanding results, CNNs have well-known shortcomings: they require extensive data and computational resources for training, they have limited theoretical guarantees and, when the training data is not representative, they may fail to generalize. Reconstruction methods, on the other hand, typically super-resolve images via optimization algorithms,

which confers them more interpretability and equips them with theoretical guarantees. They use regularizers to directly encode prior knowledge about images, e.g., that they have relatively few edges, sparse representations in some domain, or many recurring patterns. Although this direct encoding of prior knowledge dispenses reconstruction-based methods from a training stage, they are outperformed by learning-based methods.

Our framework. We propose a framework that joins reconstruction and learning techniques. Our motivating observation is that the best learning techniques, namely CNN architectures, ignore important information at test time: they fail to guarantee that the given input image matches their output when downsampled. Reconstruction algorithms, on the other hand, impose this constraint naturally. Our main idea is then to integrate the HR output of a learning-based method into a reconstruction scheme, i.e., we use an image super-resolved by a learning-based method as side information in a reconstruction-based method. For the latter, we adopt a SR algorithm based on total variation (TV) minimization, a prior that encodes the fact that images have a small number of edges. This simple prior is enough to illustrate the benefits of our framework which, nevertheless, can be adapted to more complex priors.

Our framework introduces a simple post-processing step that requires some additional computation. However, as our experiments show, it systematically improves the outputs of several CNN architectures, producing images with better peak signal to noise ratio (PSNR) and structural similarity (SSIM) [58].

2 Related Work

Single-image SR algorithms can be split into interpolation, reconstruction, and learning methods. We focus our overview on reconstruction and learning methods, as both outperform interpolation-based schemes.

2.1 Reconstruction-based SR

Reconstruction-based methods super-resolve images with an algorithm that solves an optimization problem usually containing two terms: one ensuring that the downsampled solution coincides with the input LR image and another encoding prior knowledge about images.

Statistical models. The work in [12] models the gradient profile of an image as a Gauss-Markov random field and performs SR via maximum likelihood estimation with the constraint that downsampled solution coincides with the given LR image. In [57], images are super-resolved by ensuring that the gradient of the HR image is close to a gradient profile obtained from the LR image using a similar, but parametric model, while imposing again that the downsampled solution coincides with the given LR image.

Total variation methods. A different line of work directly encodes prior knowledge into the optimization problem. For instance, [52] introduced the concept of TV to capture the number of edges in an image and observed that natural images have small TV. Since then, several algorithms have been proposed to super-resolve images by minimizing TV. For example, [24] addresses the problem by discretizing a differential equation that relates changes in the values of individual pixels to the curvature of level sets. More recently, with the development of nondifferentiable convex optimization methods, TV minimization has been addressed in the discrete setting by a wide range of algorithms [2, 9, 8, 13, 23]. There are essentially two versions of discrete TV, according to the norm that is applied to the discretized gradient at each pixel: isotropic (ℓ_2 -norm) and anisotropic (ℓ_1 -norm); see, e.g.,

[9, 23]. Since both versions are convex, discrete TV minimization algorithms not only have small computational complexity (as matrix-vector operations can be performed via the FFT algorithm), but are also guaranteed to find a global minimizer.

Alternative regularizers. Natural images also have simple representations in other domains, and other regularizers have been used for SR. For example, [52] uses the fact that images have sparse wavelet representations. The patches of natural images also tend to lie on a low-dimensional manifold [22], which has motivated nonlocal methods. These use as regularizers (nonconvex) functions that enforce many patches of an image to have similar features; see, e.g., [29, 30]. Finally, different regularizers can be combined as in [24, 35].

2.2 Learning-based SR

Learning-based methods usually operate on patches and learn how to map a LR to a HR patch. They consist of a training stage, in which the map is learned, and a testing stage, in which the map is applied to the LR patches to obtain HR patches.

Coding and dictionary learning. Coding and dictionary learning algorithms super-resolve the patches of an image by linearly combining HR patches collected (or created) from other images during training. An early method is [9], which adapts the locally linear embedding algorithm of [51] for SR. Drawing from developments in sparse-based reconstruction, [80] proposed using training images to learn two dictionaries, one for HR and another for LR patches, such that corresponding HR and LR patches have the same coefficients; see [41] for a similar approach. While these schemes rely on other images to learn the LR-HR map, [17, 36] use self-similarity to learn the map without using other images.

Regression methods. Regression methods compute the LR-HR map using a set of basis functions whose coefficients are computed via regression. The key idea is to perform clustering on the training images, and use algorithms like kernel ridge regression (KRR) [20], Gaussian process regression [16] and random forests [53] to perform SR.

CNN-based. Currently, the best performing SR methods are based on CNNs. Taking inspiration in the dictionary learning algorithm in [40], SRCNN [10] was one of the first CNNs performing image SR. It consists of a patch extraction layer, a representation layer of the non-linear mappings, and a final layer that outputs the reconstructed image. As most neural network architectures, SRCNN requires an extensive training set, and has to be retrained for each different upscaling factor. To overcome the latter problem, DRCN [19] uses a recursive convolutional layer that is repeatedly applied to obtain SR. Building on generative adversarial networks [24], state-of-the-art SR results are obtained in [2]. It proposes two networks: SRGAN, which produces photo-realistic images with high perceptual quality, and SRResNet, which achieves the best PSNR and SSIM compared to all prior SR approaches. Although CNN-based methods achieve state-of-the-art SR results, they suffer from a major shortcoming that leaves room for improvement: in general, they fail to impose consistency between the HR and LR image in the testing stage. Such consistency, however, is always enforced by reconstruction-based (specifically, optimization-based) algorithms. This observation is the main motivation behind our framework, which we introduce next.

3 Our Framework

Given a low-resolution (LR) image, our goal is to build a high-resolution (HR) version of it. Before describing how we achieve this, we introduce our model and assumptions.

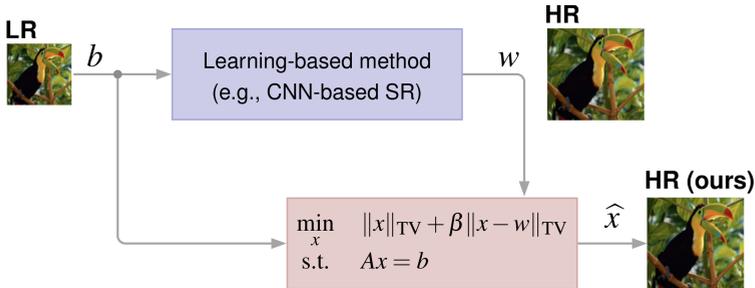


Figure 1: Our framework: a learning-based method super-resolves b into w , and both images are then used to find the final estimate \hat{x} of a TV-TV minimization problem.

Model and assumptions. We denote by $X^* \in \mathbb{R}^{M \times N}$ the HR image, and by $x^* := \text{vec}(X^*) \in \mathbb{R}^n$ its column-major vectorization, where $n := M \cdot N$. Since these represent 2D quantities, we will always work either with grayscale images or with specific channels of a given representation of color images. We assume that X^* has a small number of edges. Namely, we assume $\|x^*\|_{\text{TV}} := \sum_{i=1}^M \sum_{j=1}^N |D_{ij}^v x^*| + |D_{ij}^h x^*|$ is small, where D_{ij}^v (resp. D_{ij}^h) is a row-vector that extracts the vertical (resp. horizontal) difference at pixel (i, j) of X^* . Notice that this corresponds to the anisotropic TV-norm [9, 23], which can be written more succinctly as $\|x^*\|_{\text{TV}} = \|Dx^*\|_1$, where the rows of $D \in \mathbb{R}^{2n \times n}$ contain all the D_{ij}^v 's and D_{ij}^h 's. We assume periodic boundary conditions, so that matrix-vector products can be efficiently computed via the FFT algorithm.

The vectorization of the given LR image will be denoted by $b \in \mathbb{R}^m$, where $m < n$. We assume the LR and HR vectorizations are linearly related as $b = Ax^*$, where $A \in \mathbb{R}^{m \times n}$ is a subsampling matrix, i.e., it contains m equally interspaced rows of the identity matrix $I_n \in \mathbb{R}^{n \times n}$. Although this assumption is too simple to model certain real-case scenarios, our methodology can be immediately extended to more complex cases, e.g., when A is the product of a subsampling matrix and a blurring matrix [10, 30]. Yet, this simple assumption is sufficient to illustrate the gains that can be obtained by using our framework, explained next.

Proposed framework. Fig. 1 shows the main components of our framework. It builds on a *base method* which, because of their current state-of-the-art performance, will be a learning-based scheme. The LR image b is fed into the base method, which super-resolves it into a HR image that we will denote by $w \in \mathbb{R}^n$. As mentioned in Section 2, learning-based schemes, namely the ones based on CNN architectures, fail to enforce consistency between the LR and HR images during testing. That is, in our notation, fail to guarantee $Ax^* = b$. As shown in Fig. 1, we thus propose an additional stage, which processes both the given LR image b and the HR image w from the learning-based method to build our final estimate \hat{x} of x^* . That stage solves a problem that we call TV-TV minimization.

TV-TV minimization. Inspired by the success of sparse reconstruction schemes that integrate prior information in the form of past, similar examples [10, 28, 39], we propose to estimate x^* by solving an optimization problem whose objective balances the TV-norm of the optimization variable x and the TV-norm of the difference between x and w , while constraining $Ax = b$. That is, we obtain the final estimate \hat{x} by solving *TV-TV minimization*:

$$\begin{aligned} & \underset{x}{\text{minimize}} && \|x\|_{\text{TV}} + \beta \|x - w\|_{\text{TV}} \\ & \text{subject to} && Ax = b, \end{aligned} \quad (1)$$

where $\beta \geq 0$ is a tradeoff parameter. Since the constraints of (1) are linear and the TV-norm is convex, problem (1) is convex. To our knowledge, the first instance of (1) appeared in [10] in the context of dynamic CT imaging, where w was a coarse estimation of the image to be reconstructed. The work in [9] generalized that approach, in the context of MRI imaging, to the case where the TV-norms are weighted, and their weights are updated as in [9]. There, w represented a reference image, i.e., an image similar to the image to reconstruct.

Problem (1) enforces consistency between the LR image b and the final estimate \hat{x} via the constraint $Ax = b$. Among the possible solutions of $Ax = b$, it seeks the one that has a small TV-norm and, at the same time, does not deviate much from w , in the sense that their discrete gradients are similar. While it is also possible to enforce pixel-level similarity, we found that imposing similarity of gradients produces the best results.

As mentioned in Section 2.1, it is possible to use alternative regularizers that capture other properties of natural images, or even combine several different regularizers. Indeed, the reason why we are proposing a “framework” and not just a “method” is because we allow both blocks in Fig. 1 to vary. Using TV-TV minimization (1), however, enables us to draw from the insights in [28] for ℓ_1 - ℓ_1 minimization.

Connections with ℓ_1 - ℓ_1 minimization. Since $\|x\|_{\text{TV}} = \|Dx\|_1$, by introducing an additional variable $\mathbb{R}^{2n} \ni z = Dx$, we rewrite (1) as

$$\begin{aligned} & \underset{z,x}{\text{minimize}} && \|z\|_1 + \beta \|z - \bar{w}\|_1 \\ & \text{subject to} && Ax = b, \quad Dx = z, \end{aligned} \quad (2)$$

where we defined $\bar{w} := Dw$. We denote the concatenation of the optimization variables as $\bar{x} := (z, x) \in \mathbb{R}^{3n}$ and define $\bar{A} := \begin{bmatrix} 0_{m \times 2n} & A \\ -I_{2n} & D \end{bmatrix}$, and $\bar{b} := \begin{bmatrix} b & 0_{2n} \end{bmatrix}^\top$, where $0_{a \times b}$ (resp. 0_a) stands for the zero matrix (resp. vector) of dimensions $a \times b$ (resp. $a \times 1$), and I_{2n} is the identity matrix in \mathbb{R}^{2n} . This enables us to further rewrite (2) as

$$\begin{aligned} & \underset{\bar{x}}{\text{minimize}} && \|P_{2n}\bar{x}\|_1 + \beta \|P_{2n}\bar{x} - \bar{w}\|_1 \\ & \text{subject to} && \bar{A}\bar{x} = \bar{b}, \end{aligned} \quad (3)$$

where $P_{2n} \in \mathbb{R}^{2n \times 3n}$ contains the first $2n$ rows of the identity matrix in \mathbb{R}^{3n} , I_{3n} . That is, for a vector $v \in \mathbb{R}^{3n}$, $P_{2n}v$ extracts the first $2n$ components of v . The work in [28] characterized the performance of the problem in (3) when P_{2n} is the full identity matrix, rather than just part of it, a problem they called ℓ_1 - ℓ_1 minimization. Assuming that \bar{A} has Gaussian random entries, [28] provided precise reconstruction guarantees for such a problem. In particular, they showed that $\beta = 1$ yields the best performance in theory and in practice, and also suggested a heuristic that improves the performance of the method: amplifying w by a small factor, i.e., using $\alpha\bar{w}$ instead of \bar{w} in (3), for an α slightly larger than 1.

Although these results were proved only for the case in which \bar{A} is a Gaussian matrix, experiments in [27] suggest that they also hold for other types of matrices. This led us to select $\beta = 1$ and $\alpha = 1.2$ in all our experiments.

Solving (1). A disadvantage of our framework compared to learning-based methods, in particular CNN architectures, is that solving an optimization problem like (1) requires some computation. Yet, because the subsampling matrix A and difference matrix D are very structured, we can design algorithms that take advantage of fast matrix-vector multiplications. Specifically, given $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^m$, Au outputs a subvector of u , and $A^\top v$ outputs an n -dimensional vector whose entries are the entries of v (at locations specified by A) or zeros.

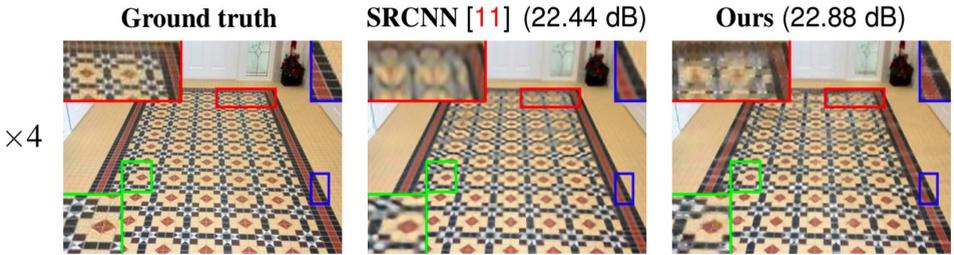


Figure 2: Performance of our method, using SRCNN [11] as the base method, for an image not included in any of the training or test sets. The ground truth was downsampled by 4 and then super-resolved with SRCNN and our method. PSNR values are shown in parenthesis.

These operations require neither explicit construction of A , nor any floating-point operation (just memory access). Because we assume periodic boundary conditions, the matrix-vector products Du and $D^T v$, for $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^{2n}$, can be efficiently computed using the FFT algorithm, which requires $O(n \log n)$ floating-point operations; see [23].

To solve (1), we apply ADMM [9], guaranteeing that its implementation uses only matrix-vector products as above. ADMM is applied directly to problem (2) with the splitting $f(z, x) = \|z\|_1 + \beta \|z - w\|_1$ and $g(x, z) = \mathbf{i}_{\{(z, x): Ax=b, Dx=z\}}(z, x)$, where $\mathbf{i}_S(u)$ is an indicator function, i.e., $\mathbf{i}_S(u) = 0$ if $u \in S$, and $\mathbf{i}_S(u) = +\infty$ otherwise. Our focus was not on obtaining the most efficient implementation, but we point out possible improvements in Section 5.

4 Experiments

We tested our framework on the standard test sets Set5 [8], Set14 [40], and BSD100 [25], each containing respectively 5, 14, and 100 images. For RGB images, we only super-resolved the luminance channel of the YCbCr color space, as in [10, 17, 19, 21]. All our experiments were run on the same machine with 8GB of RAM and a 3.1 GHz Intel i5 CPU.

We used five state-of-the-art learning schemes: KRR by Kim [20], self-similarity [17], pre-trained CNNs [10, 19, 21] and basic TV minimization, i.e., (1) with $\beta = 0$, using TVAL3 [23]. Of these methods, only [23] operates on full images, just like our TV-TV minimization algorithm; all the remaining methods operate only on patches. The output images of SRCNN [10], Kim [20] and SelfExSR [17] were obtained from the repository¹.

As our framework is tied to a given base method, whenever the latter changes, so does our method. For assessing performance, we used PSNR and SSIM [38] computed on the luminance channel. We mention that [20] proposes two algorithms, of which we use the SRResNet, since it yields the best PSNR/SSIM values.

4.1 Illustrating consistency between LR and HR images

To demonstrate how imposing consistency between the LR and HR images can improve SR performance, we conducted an experiment using an image that is not in any of the above datasets. The image, of dimensions 472×352 , is shown in the left of Fig. 2 and was selected because of its structured patterns. We downsampled it by a factor of 4, and then reconstructed it with the CNN architecture SRCNN [10], and with our method using SRCNN as the base

¹<https://github.com/jbhuang0604/SelfExSR>

algorithm. Since SRCNN fails to enforce consistency between the LR and HR images during testing, its output (Fig. 2, center) blurs the tiles and removes some details. In contrast, our method, which uses the output of SRCNN (Fig. 2, center) and the $4\times$ downsampled version of the ground truth (Fig. 2), preserves details of the tiles and reduces blurring. The improvement in image quality, however, comes with a computational cost: to super-resolve that image required SRCNN 2 seconds, and solving (1) required 557 seconds.

4.2 Experiments on Set4, Set5, and BSD100

Tab. 1 shows the results of our experiments for the testing datasets Set5, Set14, and BSD100. It displays the PSNR and SSIM values for $2\times$ and $4\times$ upscaling tasks, and these values represent averages over all the images in the corresponding dataset. Recall that higher PSNR and SSIM values reflect better performance. The table has 5 parts, one for each variation of the base method. For example, the first 3 rows refer to the results when the method in Kim [20] is the base algorithm; in the next 3 rows, the base method is SRCNN [10], and so on. The values for TVAL3 [23] are repeated in all parts of the table, since its input is the same image in all the cases. We note that table displays results for SRResNet [21] only for $4\times$ upscaling, as it was designed specifically for that factor.

It can be observed that TVAL3, an algorithm for solving (1) with $\beta = 0$, yields the worst performance. All the learning-based methods (4th column) systematically give results better than TVAL3. Our method (last column), in turn, uses the outputs of those algorithms and, by enforcing consistency between the LR and HR images, systematically improves their outputs: for example, the range of the gains in PSNR values was 0.65-2.29 dB.

Fig. 3 shows the reconstructions of a particular image in Set14 for all the methods. The top-left box corresponds to the ground truth and TVAL3, which were the same in all the experiments. The remaining boxes show the results of a given learning-based method and the corresponding output using our framework. As in Section 4.1, it can be seen that the learning-based methods tend to produce images that lack details, as they fail to impose consistency with the given LR image. Our method, in contrast, preserves fine details and does not introduce as much blurring. The execution times of SRCNN [10] and DRCN [19] to super-resolve the sample image in Fig. 3 are shown in Tab. 2. Although our method yields better PSNR/SSIM values, it requires significantly more computational time.

5 Conclusions

We proposed a framework for super-resolving images that leverages the good performance of learning-based methods, including CNN architectures, and the consistency enforced by reconstruction-based algorithms. The framework takes a base learning-method, and improves its output by solving a TV-TV minimization problem. Although the image model we used is quite simple, it suffices to illustrate the reconstruction gains in PSNR and SSIM that can be obtained. Our experimental results involving several state-of-the-art SR algorithms show that the proposed framework achieves significant performance gains over those algorithms, however, at an additional computation cost.

Possible research directions include using regularizers that capture more realistic image assumptions and accelerating the algorithm for solving the TV-TV minimization algorithm e.g., using a C implementation, or a trained recurrent neural network as a solver [15].

Table 1: Average PSNR (SSIM) results using the reference methods.

Dataset	Scale	TVAL3 [23]	Kim [24]	Ours
Set5	×2	34.23 (0.95)	36.24 (0.95)	38.40 (0.97)
	×4	26.87 (0.83)	30.07 (0.86)	31.35 (0.90)
Set14	×2	30.83 (0.91)	32.14 (0.90)	34.15 (0.94)
	×4	24.58 (0.74)	27.18 (0.74)	28.09 (0.82)
BSD100	×2	30.10 (0.89)	31.11 (0.88)	32.80 (0.93)
	×4	24.62 (0.71)	26.71 (0.70)	27.46 (0.79)
Dataset	Scale	TVAL3 [23]	SRCNN [14]	Ours
Set5	×2	34.23 (0.95)	36.28 (0.95)	38.48 (0.97)
	×4	26.87 (0.83)	30.08 (0.85)	31.38 (0.90)
Set14	×2	30.83 (0.91)	32.00 (0.90)	34.17 (0.94)
	×4	24.58 (0.74)	27.13 (0.74)	28.13 (0.82)
BSD100	×2	30.10 (0.89)	31.11 (0.89)	32.83 (0.93)
	×4	24.62 (0.71)	26.70 (0.70)	27.46 (0.79)
Dataset	Scale	TVAL3 [23]	SelfExSR [17]	Ours
Set5	×2	34.23 (0.95)	36.50 (0.95)	38.79 (0.97)
	×4	26.87 (0.83)	30.33 (0.86)	31.62 (0.91)
Set14	×2	30.83 (0.91)	32.23 (0.91)	34.36 (0.94)
	×4	24.58 (0.74)	27.40 (0.75)	28.34 (0.82)
BSD100	×2	30.10 (0.89)	31.18 (0.89)	32.91 (0.93)
	×4	24.62 (0.71)	26.85 (0.71)	27.61 (0.79)
Dataset	Scale	TVAL3 [23]	DRCN [19]	Ours
Set5	×2	34.23 (0.95)	37.63 (0.96)	39.47 (0.97)
	×4	26.87 (0.83)	31.53 (0.89)	32.59 (0.83)
Set14	×2	30.83 (0.91)	33.06 (0.91)	34.92 (0.95)
	×4	24.58 (0.74)	28.03 (0.77)	28.77 (0.83)
BSD100	×2	30.10 (0.89)	31.85 (0.89)	33.40 (0.93)
	×4	24.62 (0.71)	27.24 (0.72)	27.90 (0.80)
Dataset	Scale	TVAL3 [23]	SRResNet [21]	Ours
Set5	×4	26.87 (0.83)	32.06 (0.89)	32.98 (0.92)
Set14	×4	24.58 (0.74)	28.59 (0.78)	29.24 (0.84)
BSD100	×4	24.62 (0.71)	27.60 (0.74)	28.19 (0.81)



Figure 3: Reconstructions of the Monarch image (Set14) for all the methods in Tab. 1. Some details are highlighted (best viewed by zooming in on a computer screen).

Table 2: Execution times (in secs) for reconstructing the image in Fig. 3.

Scale	TVAL3 [23]	Base: SRCNN [11]		Base: DRCN [19]	
		SRCNN	Ours	DRCN	Ours
×2	3.34	9.31	411.46	156.20	729.08
×4	6.08	9.43	770.56	143.20	1456.04

6 Acknowledgments

Work supported by EPSRC EP/S000631/1 via UDRC/MoD.

References

- [1] H. A. Aly and E. Dubois. Image up-sampling using total-variation regularization with a new observation model. *IEEE Transactions on Image Processing*, 14(10):1647–1659, 2005.
- [2] S. Becker, J. Bobin, and E. Candès. NESTA: a fast and accurate first-order method for sparse recovery. *SIAM J. Imaging Sciences*, 4(1):1–39, 2011.
- [3] M. Bevilacqua, A. Roumy, C. Guillemot, and M. L. A. Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *British Machine Vision Conference (BMVC)*, pages 135.1–135.10, 2012.
- [4] J. M. Bioucas-Dias and M. A. T. Figueiredo. A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration. *IEEE Transactions on Image Processing*, 16(12):2992–3004, 2007.
- [5] T. Blu, P. Thevenaz, and M. Unser. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5):710–719, 2004.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [7] E. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14:877–905, 2008.
- [8] A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20(1):89–97, 2004.
- [9] H. Chang, D. Yueng, and Y. Xiong. Super-resolution through neighbor embedding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 275–282, 2004.
- [10] G.-H. Chen, J. Tang, and S. Leng. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly under-sampled projection data sets. *Medical physics*, 35(2):660–663, 2008.
- [11] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 184–199, 2014.
- [12] R. Fattal. Image upsampling via imposed edge statistics. *ACM SIGGRAPH*, 26(3), 2007.
- [13] T. Goldstein and S. Osher. The split Bregman method for ℓ_1 -regularized problems. *SIAM J. Imaging Science*, 2(2):323–343, 2009.

- [14] I. Goodfellow, J. P. Abadie, M. Mirza, B. Xu, D. W. Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Neural Information Processing Systems (NIPS)*, pages 2672–2680, 2014.
- [15] K. Gregor and Y. LeCun. Learning fast approximations of sparse coding. In *International Conference on Machine Learning (ICML)*, pages 399–406, 2010.
- [16] H. He and W. Siu. Single image super-resolution using gaussian process regression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 449–456, 2011.
- [17] J. Huang, A. Singh, and N. Ahuja. Single image super-resolution from transformed self-exemplars. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5197–5206, 2015.
- [18] R. G. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-29(6):1153–1160, 1981.
- [19] J. Kim, J.K. Lee, and K.M. Lee. Deeply-recursive convolutional network for image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1637–1645, 2016.
- [20] K. I. Kim and Y. Kwon. Single-image super-resolution using sparse regression and natural image prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1127–1133, 2010.
- [21] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114, 2017.
- [22] A. B. Lee, K. S. Pedersen, and D. Mumford. The nonlinear statistics of high-contrast patches in natural images. *International Journal of Computer Vision*, 54(1-3):83–103, 2003.
- [23] C. Li, W. Yin, H. Jiang, and Y. Zhang. An efficient augmented lagrangian method with applications to total variation minimization. *Computational Optimization and Applications*, 56(3):507–530, 2013.
- [24] J. Li, J. Wu, H. Deng, and J. Liu. A self-learning image super-resolution method via sparse representation and non-local similarity. *Neurocomputing*, 184:196–206, 2016.
- [25] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *International Conference on Computer Vision (ICCV)*, volume 2, pages 416–423, 2001.
- [26] B. S. Morse and D. Schwartzwald. Image magnification using level set reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 333–334, 2001.

- [27] J. F. C. Mota, N. Deligiannis, A. C. Sankaranarayanan, V. Cevher, and M. R. D. Rodrigues. Adaptive-rate reconstruction of time-varying signals with application in compressive foreground extraction. *IEEE Transactions on Signal Processing*, 64(14):3651–3666, 2016.
- [28] J. F. C. Mota, N. Deligiannis, and M. R. D. Rodrigues. Compressed sensing with prior information: Strategies, geometry, and bounds. *IEEE Transactions on Information Theory*, 63(7):4472–4496, 2017.
- [29] G. Peyré, S. Bougleux, and L. Cohen. Non-local regularization of inverse problems. In *European Conference on Computer Vision (ECCV)*, pages 57–68, 2008.
- [30] M. Protter, M. Elad, H. Takeda, and P. Milanfar. Generalizing the nonlocal-means to super-resolution reconstruction. *IEEE Transactions on Image Processing*, 18(1):36–51, 2009.
- [31] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.
- [32] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
- [33] S. Schuler, C. Leistner, and H. Bischof. Fast and accurate image upscaling with super-resolution forests. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3791–3799, 2015.
- [34] P. Sen and S. Darabi. Compressive image super-resolution. In *Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers*, pages 1235–1242, 2009.
- [35] F. Shi, J. Cheng, L. Wang, P.-T. Yap, and D. Shen. LRTV: MR image super-resolution with low-rank and total variation regularizations. *IEEE Transactions on Medical Imaging*, 34(12):2459–2466, 2015.
- [36] A. Singh, F. Porikli, and N. Ahuja. Super-resolving noisy images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2846–2853, 2014.
- [37] J. Sun, Z. Xu, and H. Shumand. Image super-resolution using gradient profile prior. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.
- [38] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [39] L. Weizman, Y. C. Eldar, and D. Ben-Bashat. Reference-based MRI. *Med. Phys.*, 43(10):5357–5369, 2016.
- [40] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.
- [41] R. Zeyde, M. Elad, and M. Protter. On single image scale-up using sparse representations. In *Curves and Surfaces*, pages 711–730, 2012.