

# Super-resolution via Transform-invariant Group-sparse Regularization

Carlos Fernandez-Granda  
Stanford University  
cfgranda@stanford.edu

Emmanuel J. Candès  
Stanford University  
candes@stanford.edu

## Abstract

We present a framework to super-resolve planar regions found in urban scenes and other man-made environments by taking into account their 3D geometry. Such regions have highly structured straight edges, but this prior is challenging to exploit due to deformations induced by the projection onto the imaging plane. Our method factors out such deformations by using recently developed tools based on convex optimization to learn a transform that maps the image to a domain where its gradient has a simple group-sparse structure. This allows to obtain a novel convex regularizer that enforces global consistency constraints between the edges of the image. Computational experiments with real images show that this data-driven approach to the design of regularizers promoting transform-invariant group sparsity is very effective at high super-resolution factors. We view our approach as complementary to most recent super-resolution methods, which tend to focus on hallucinating high-frequency textures.

## 1. Introduction

A fundamental challenge in computer vision and image processing is to increase the resolution of blurry images. Multi-frame super-resolution consists of combining several low-resolution images for this purpose [6]. Single-frame super-resolution is even more challenging. The aim is to obtain a higher-resolution image by upsampling a single image. In general, it is of course impossible to recover fine scale details that are absent from the low-resolution image. We can only hope to reconstruct certain very specific structures (see [4] for theoretical results on the super-resolution of pointwise objects) or to *hallucinate* high-frequency textures that are visually pleasing [1].

In this work we consider planar regions taken from 3D scenes that have straight edges aligned in a few main directions, such as the one in Figure 1. This class of surfaces is of great practical interest. They are ubiquitous in urban environments and recent large-scale urban 3D mapping efforts (such as the Apple 3D map) make such data readily avail-




Figure 1. Image of a building façade (top left) and its gradient (top right) along with their low-rank representation (bottom).

able. Existing super-resolution techniques can be applied to this class of textures to obtain reasonably good upsampling results up to factors of three or four. However, most of these techniques are designed for generic images and do not take into account what high-level structure may be present in the data. In this work, we explore the possibility of attaining higher upsampling factors by harnessing such prior knowledge for images with structured edges.

Unfortunately, the non-uniform blur and deformations induced by the projection of 3D surfaces onto the imaging plane make it very challenging to exploit prior knowledge about the structure of the data directly. In fact, in our quest to super-resolve regions with structured edges we face two fundamental questions that are at the core of many problems in computer vision:

1. How should we leverage prior knowledge about global features of the 3D geometry of a scene?
2. How can we deal with the camera projection, which distorts these global features?

In our case, the global features include straight edges oriented in a few main directions, which suggests promoting structured sparsity of the image gradient, for example via



Figure 2. Eight-fold super-resolved reconstruction obtained by our method (bottom right) of an approximately low-rank region (top right and white rectangle on left) in a low-resolution image of a truck (left).

a group  $\ell_1$ -norm penalty [18]. This provides a partial answer to the first question. However, we must still tackle the second. Indeed, if we only have access to a projected low-resolution image, it is unclear how to design an appropriate group-sparsity penalty for the image gradient. A solution in the case of highly structured 3D scenes is to use recent advances in the recovery of *low-rank textures* [19], defined as low-rank structures deformed by affine or projective transformations. These techniques allow to obtain a transformation that reveals low-rank structure from the low-resolution image. In the case of images with highly structured edges, the sparsity pattern of the gradient tends to follow a low-rank pattern, as illustrated by Figure 1. Our method combines these insights by:

1. Learning the domain transform that reveals the low-rank structure of the data.
2. Designing a nonparametric regularizer tailored to this structure.

Under the assumptions that the edge structure of the 3D scene is approximately low rank, this data-driven procedure produces a convex regularizer that allows to super-resolve the image very effectively. Figure 2 illustrates the power of our approach with a simple example in which it deconvolves lines that are almost merged together due to heavy aliasing in the low-resolution data.

Our approach departs significantly from current state-of-the-art methods for single-image super-resolution, which infer the fine-scale details of the upsampled image from a database of corresponding low and high-resolution patches. Different options to perform this inference include fitting a Markov random field model [8], learning a parametric model for the edge profiles [13], or applying a sparse-coding framework [17]. These *example-based* methods enforce local consistency to produce sharp-looking edges and

are able to hallucinate high-frequency textures very effectively at moderate upsampling factors, especially if prior knowledge about these textures is available [9, 7, 14]. However, they are not well adapted to deal with global features, such as the straight edges in Figure 1. Our work is designed to exploit such features and is consequently complementary to these methods. Merging both approaches is an interesting direction for future research.

To recapitulate, our main contribution is a principled methodology for the super-resolution of planar regions with regular structures, which achieves high-quality results at upsampling factors that are problematic for other methods. This is achieved by learning the geometric structure of the 3D scene and leveraging it within a variational framework. We motivate and explain our technique in more detail in Section 2. Experimental results are provided in Section 3. In Section 4 we end the paper with some conclusions and ideas for further work.

## 2. Proposed method

### 2.1. Directional total variation

Consider the problem of designing a regularizer adapted to the problem of super-resolving images with sharp edges oriented in a few main directions. The most common choice in the literature is to penalize the  $\ell_1$  norm of the gradient [11], also known as the total variation (TV) of the image, or related non-convex penalties [10] in order to obtain an estimate with a sparse gradient (see also [15] for a recent approach that takes discretization into account). Unfortunately, minimizing the total variation often fails to super-resolve two-dimensional edges, even in the case of very simple piecewise-constant images such as the checkerboard shown in Figure 3. This failure is largely due to the fact that the regularizer is agnostic to the orientation of the edges in the image, and in particular to the correlation between the orientation of nearby edges. This suggests resorting to a regularizer that is better adapted to the high-level structure of the image gradient.

Let us assume that, as is the case for the checkerboard in Figure 3, we happen to know the directions of most edges in the image. In this case, the gradient in the image is not only sparse, but *group sparse* [18], since its nonzero elements are grouped along horizontal and vertical lines. As a result, a more suitable regularizer is the directional total variation (DTV) of the image, defined as

$$\text{DTV}(I) = \sum_{x=1}^{N_1} \sqrt{\sum_{y=1}^{N_2-1} (I(x, y+1) - I(x, y))^2} + \sum_{y=1}^{N_2} \sqrt{\sum_{x=1}^{N_1-1} (I(x+1, y) - I(x, y))^2}, \quad (2.1)$$

for an image  $I \in \mathbb{R}^{N_1 \times N_2}$ . In words, this cost function is equal to the sum of the  $\ell_2$  norms of the difference between adjacent rows and columns. It is designed to favor edges that are aligned horizontally and vertically. A similar regularizer has been proposed for multiple change-point detection in time-series analysis [3]. The top of Figure 3 compares the results of minimizing the TV and DTV cost functions to perform nonblind deblurring of a checkerboard image. DTV minimization recovers the edge structure of the image very precisely, whereas TV minimization produces significant artifacts.




Figure 3. Nonblind deblurring of a checkerboard using TV and DTV minimization. The blurring kernel was a  $21 \times 21$  Gaussian kernel with standard deviation equal to 3.

### 2.2. Transform-invariant regularization

The DTV regularizer proposed in Section 2.1 is obviously of very limited applicability. We seldom encounter images where the edges are perfectly aligned horizontally and vertically. Because of this, we turn towards a more general model, which applies to many man-made structures such as building façades. We consider images such that there exists an affine or projective transform  $\tau$  for which most of the edges of  $I \circ \tau$  are aligned vertically and horizontally. To be clear,  $I \circ \tau$  is a new image obtained by applying  $\tau$  to the domain of  $I$  and then resampling. In general it is not straightforward to design a regularizer adapted to such a model. The reason is that the gradient is no longer group sparse along a few main directions. However, *it is group sparse modulo the transform  $\tau$* . Following this insight, we define

$$\text{TI-DTV}(I) = \text{DTV}(I \circ \tau), \quad (2.2)$$

a cost function that promotes straight edges in the transformed image, where TI-DTV stands for transform-invariant directional total variation. Indeed, the cost function (2.2) is invariant to affine or projective transforms of  $I$ , as long as we are able to estimate them a priori. As we will see, this allows to factor out significant deformations induced by the camera projection. At the bottom of Figure 4 we can see the results of minimizing the TI-DTV cost function to perform nonblind deblurring of a tilted checkerboard. Unlike TV minimization, TI-DTV minimization yields a sharp and aliasing-free result, almost identical to the original.




Figure 4. Nonblind deblurring of a tilted checkerboard using TV and TI-DTV minimization. The blurring kernel was a  $21 \times 21$  Gaussian kernel with standard deviation equal to 3.

### 2.3. Transform-invariant low-rank textures

In order to use the cost function proposed in (2.2), it is necessary to learn a transform “ $\tau$ ” mapping the image to a domain where its edges are mostly aligned vertically and horizontally. We propose doing this by exploiting the fact that images with vertical and horizontal edges tend to be approximately low rank when viewed as a matrix. This is obviously the case for the checkerboard in Figure 4, but holds much more broadly. For example, Figure 5 shows a low-rank texture extracted from a blurry image. The main edges are indeed aligned horizontally and vertically by the transformation associated to the low-rank texture. An image with low-rank structured edges might lose its sharpness at low resolutions, but it remains an approximately low-rank texture. This is crucial for our interests, since we can consequently use TILT to learn the transform associated to the edge structure and then apply the regularizer proposed in Section 2.2.




Figure 5. Blurry image containing an approximately low-rank texture (left) together with the result of rectifying it using the transform learnt by TILT (right).

The authors of [19] develop robust computational tools that allow to extract low-rank textures distorted by affine or projective transformations. With a slight abuse of terminology, we refer to this method as TILT, which stands for transform-invariant low-rank textures. In essence, TILT al-

lows to compute a transform  $\tau$  from an image  $I$  such that  $I \circ \tau = L + E$ , where  $L$  is a low-rank matrix and  $E$  accounts for sparse deviations from the low-rank model. This is achieved by solving

$$\min_{\tau, L, E} \|L\|_* + \lambda \|E\|_1 \quad \text{subject to } I \circ \tau = L + E, \quad (2.3)$$

which is inspired by theoretical results on the recovery of sparsely corrupted low-rank matrices [5]. Even though Problem (2.3) is non-convex, it can often be tackled efficiently by solving a sequence of convex programs. We refer to [19] for more details on transform-invariant low-rank textures and on how to solve Problem (2.3), but we would like to mention that the presence of the sparse term  $E$  is of vital importance if we apply TILT to retrieve low-rank textures from low-resolution inputs. The reason is that it accounts for artifacts caused by blur and pixelation. This is illustrated by Figure 6, which shows the low-rank and sparse components obtained from a blurry image.




Figure 6. Low-rank (center) and sparse (right) components extracted by TILT from the low-resolution image on the left.

## 2.4. Super-resolution via TI-DTV regularization

We finally have all the tools to tackle the problem of super-resolving an image obtained from a 3D scene with structured edges. We propose to leverage a data-driven convex regularizer adapted to the 3D geometry revealed by TILT. Given a low-resolution image  $I_{LR} \in \mathbb{R}^{n_1 \times n_2}$ , we aim to obtain a higher-resolution image  $I_{HR} \in \mathbb{R}^{N_1 \times N_2}$ ,  $N_1 > n_1$ ,  $N_2 > n_2$  such that

$$I_{LR} \approx \mathcal{D}(K \otimes I_{HR}), \quad (2.4)$$

for a downsampling operator  $\mathcal{D} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^{n_1 \times n_2}$  and a blurring kernel  $K \in \mathbb{R}^{N_1 \times N_2}$ . In order to do this, we suggest a two-step method.

1. We apply TILT to the low-resolution image  $I_{LR}$  in order to obtain a transform  $\tau$  which reveals the low-rank edge structure of the image. In practice, we upsample  $I_{LR}$  using bicubic interpolation before learning the transformation. To implement this step we use the code for TILT available online.

2. We obtain a high-resolution image by solving

$$\min_{\tilde{I}_{HR}} \left\| \mathcal{D}(K \otimes \tilde{I}_{HR}) - I_{LR} \right\|_2 + \lambda \cdot \text{DTV}(A_\tau \tilde{I}_{HR}) + \beta \cdot \text{TV}(\tilde{I}_{HR}), \quad (2.5)$$

where  $\lambda$  and  $\beta$  are regularization parameters, TV represents the usual total variation operator for discrete images (i.e. the sum of the horizontal and vertical finite differences) and  $A_\tau$  is a linear operator that maps the image to the domain where we seek to penalize the directional total variation.

For color images we apply this procedure to the illuminance channel and upsample the chrominance components Cb and Cr using bicubic interpolation.

The cost function in (2.5) combines a data fidelity term, which enforces the model (2.4), with the convex regularizer described in Section 2.2 and also with an extra term that penalizes the total variation of the image. The reason is that the TI-DTV term is not enough to stabilize the reconstruction as  $A_\tau$  is often rank deficient. Additionally, the careful reader might remark that we assume that the low pass kernel  $K$  is known. In practice, we use a Gaussian kernel with a standard deviation  $\sigma$  slightly greater than the upsampling factor divided by two. In general, the algorithm is quite robust to this choice and also to the values of  $\lambda$  and  $\beta$ . For the results in Section 3 we used  $\sigma = 5.5$  in all cases and  $\lambda = 3$  and  $\beta = 0.1$  in most cases. The Supplementary Material<sup>1</sup> elaborates on the stability of our method to changes in the parameters  $\sigma$ ,  $\lambda$  and  $\beta$ .

In order to solve Problem (2.5) we apply the Templates for First-Order Conic Solvers (TFOCS) framework proposed in [2], code for which is provided online. This framework consists in casting the problem as a conic program, determining the dual problem and applying a first-order method, such as accelerated gradient descent, to solve a smoothed version of the dual. The main advantage of TFOCS over other solvers is that it allows to minimize functions of the form  $\|Wx\|$  for an arbitrary linear operator  $W$ , a vector  $x$  and any norm for which we can compute the corresponding dual norm efficiently. We can apply TFOCS almost off the shelf to solve (2.5) by implementing functions to apply the operator  $A_\tau$  and to compute the dual of the mixed  $\ell_1/\ell_2$  norm, equal to the mixed  $\ell_\infty/\ell_2$  norm. Some care is needed in the implementation of  $A_\tau$  for the solver to run quickly. The optimization algorithm applies the adjoint  $A_\tau^*$  repeatedly, so it is important to make this efficient. This can be done by implementing  $A_\tau$  with a sparse matrix that samples the transformed image on a new grid using bilinear interpolation. We refer to [2] for further relevant details on TFOCS, in particular Sections 4.4, 4.5 and 4.6.

<sup>1</sup>[www.stanford.edu/~cfgranda/II\\_DTV\\_supp\\_mat.pdf](http://www.stanford.edu/~cfgranda/II_DTV_supp_mat.pdf)

Finally, we would like to point out that (2.5) is a convex cost function. As a result, the algorithm is robust to the choice of initialization and converges to the same solution even if we use the optimum for TV or other cost functions such as (2.6) as a starting point.

### 2.5. An alternative approach

Instead of solving Problem (2.5), another option to super-resolve our class of images of interest is to work directly in the rectified domain penalizing the DTV norm. More precisely, one can compute the rectified low-resolution image  $I_{LR}^\tau = I_{LR} \circ \tau$  and then solve the optimization problem

$$\min_{\tilde{I}_{HR}^\tau} \left\| \mathcal{D} \left( K \otimes \tilde{I}_{HR}^\tau \right) - I_{LR}^\tau \right\|_2 + \alpha \cdot \text{DTV} \left( \tilde{I}_{HR}^\tau \right), \quad (2.6)$$

where  $\alpha$  is an optimization parameter. The super-resolved image in the original domain can then be obtained from the solution by inverting  $\tau$ . This simplifies the optimization problem that we must solve, but yields worse results than the method proposed in Section 2.4. The reason is that we are implicitly assuming that the kernel acts on the rectified domain, a model which is inaccurate and does not account for the uneven blur caused by the camera projection. Adequate modeling of the downsampling operator is crucial to super-resolve effectively [12], so it not surprising that the results for this alternative method are not as sharp as those obtained with TI-DTV regularization, as shown in Figure 7.



Figure 7. Comparison between solving Problems (2.5) and (2.6) to super-resolve at an upsampling factor of 8 using geometric information obtained from the low-resolution image. TI-DTV produces superior results.

## 3. Experiments

The code used for the experiments is available at [www.stanford.edu/~cfgranda/TI\\_DTV\\_code.zip](http://www.stanford.edu/~cfgranda/TI_DTV_code.zip).

### 3.1. Super-resolution of real images

In this section we provide some experimental results to evaluate the performance of our method. We focus on qualitative comparisons, since there is no clear metric capable of quantifying the quality of super-resolved images (for instance, at high upsampling factors the mean-square error can be better for blurry images that do not enhance any features of interest). Further experiments are reported in the Supplementary Material. In our first example, we take large planar regions from five images in the SUN database [16], shown in Figure 8, and compare our method with other representative super-resolution methods developed in the literature. Although we apply the algorithms to the whole planar region, zoomed-in areas are shown due to space limitations. Interpolation algorithms are represented by bicubic interpolation, which we compute using the Matlab function *imresize*. To compare with classical variational techniques we use TFOCS to solve the optimization problem

$$\min_{\tilde{I}_{HR}} \left\| \mathcal{D} \left( K \otimes \tilde{I}_{HR} \right) - I_{LR} \right\|_2 + \gamma \text{TV} \left( \tilde{I}_{HR} \right), \quad (3.1)$$

where  $\gamma$  is a regularization parameter, since TV regularization is usually the method of choice to promote sharp edges in image processing. Finally, we choose the sparse-coding super-resolution algorithm of Yang et al [17] as a representative of exemplar-based algorithms using patches. In the literature, this method is competitive with most other exemplar-based techniques. It is actually noted in [17] that exemplar-based methods have difficulties dealing with highly structured textures such as building facades, because it is difficult to build a dictionary that can exhaust edges in all directions and scales. Nevertheless, it serves as a baseline to compare with our method. We use the code available



Figure 8. Planar regions with structured edges extracted from images in the SUN database [16]. A white rectangle highlights the areas shown in Figure 12.

online for this algorithm, which allows to apply an upsampling factor of 4. For the rest of the methods, including ours, we apply an upsampling factor of 8.

Zoomed-in areas of the results are shown in Figure 12. In all cases bicubic interpolation produces images that are very blurry. The results for TV regularization are sharper, but they contain significant artifacts which make edges appear wobbly instead of straight. Despite its reduced upsampling factor, the sparse-coding algorithm is also not capable of super-resolving edges effectively and its results are only slightly better than those of bicubic interpolation. In contrast, TI-DTV regularization produces clear straight edges that correspond to the global geometry of the planar surface, yielding upsampled images that are significantly sharper than the rest. Figure 13 shows the effect of applying the transform learnt from the low-resolution images to the high-resolution images obtained by TI-DTV regularization. As expected the edges align mostly horizontally and vertically following the low-rank structure.

For the top example in Figure 12, where the low-resolution image has size 120x136, the running time required by TI-DTV regularization is of 123.0 s (SRF=2), 417.8 s (SRF=4) and 1713.5 s (SRF=8) on a desktop computer with a 3.2 GHz processor and 4 GB of RAM.

### 3.2. Super-resolution of text

Text follows the model that we consider to some extent, since most letters contain horizontal or vertical edges. As a result, TI-DTV regularization is capable of effectively super-resolving letter or characters printed on distorted surfaces. To demonstrate this, Figure 11 shows four such examples and compares the results obtained from bicubic interpolation and our method. Comparisons to other methods are included in the Supplementary Material. TI-DTV regularization is clearly superior in all cases, despite the fact that some letters have edges that are not aligned with the low-rank structure (see the following section).

### 3.3. Limitations

As we have made clear throughout this paper, our method is geared to the super-resolution of planar surfaces that are approximately low-rank and have straight edges that are oriented following the low-rank structure. If these conditions are not met, the algorithm might produce artifacts in regions that resemble horizontal or vertical edges. In Figure 9 such artifacts can be seen along the jagged diagonal line. A way to overcome this would be to incorporate other orientations into the transform-invariant group-sparsity penalty, or introduce a more sophisticated group-sparsity prior adapted to such features. In any case, the method often degrades gracefully in regions that do not have straight edges with the right orientation. In Figure 10, for example, TI-DTV does not produce any artifacts around the arc and

at the same time super-resolves sharply the edges in the rest of the image. In comparison, TV makes the arc look sharper but generates obvious artifacts.



Figure 9. Eight-fold upsampled detail (right) of a shop sign (left) showing the kind of artifact that might be produced by TI-DTV regularization.



Figure 10. Comparison of the results of super-resolving an image that does not completely conform to the transform-invariant low-rank model using bicubic interpolation, TV regularization and TI-DTV regularization. The upsampling factor is 8.

## 4. Conclusion and extensions

We believe that developing tools capable of constraining non-local image structure is a crucial step towards achieving high-quality super-resolution at large upsampling factors. Our contributions are the introduction of a principled methodology in which such constraints are imposed through data-driven non-parametric regularizers and a robust implementation of this methodology for a particular class of images, which yields state-of-the-art results. Future research directions include combining our framework with patch-based methods, designing group-sparsity patterns adapted to other classes of images and developing alternative approaches to learn the parametric transformations that make our regularizers approximately invariant to the camera projection.

## Acknowledgements

Part of this work was carried out when C.F. was an intern at Microsoft Research Asia. C.F. would like to thank Yi Ma for valuable comments and Ewout van den Berg for useful discussions. C.F. was funded by a Fundación Caja Madrid Fellowship. E.C. was partially funded by the Simons Foundation and by a gift from the Broadcom Foundation.



Figure 11. Eight-fold upsampling of surfaces containing text taken from the SUN database [16]. The input is shown on the left with the right scaling for reference.

## References

- [1] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *Pattern Analysis and Machine Intelligence, IEEE Trans. on*, 24(9):1167–1183, 2002. 1
- [2] S. R. Becker, E. J. Candès, and M. Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical programming computation*, 3(3):165–218, 2011. 4
- [3] K. Bleakley and J.-P. Vert. The group fused lasso for multiple change-point detection. *arXiv:1106.4199*, 2011. 3
- [4] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*. To appear. 1
- [5] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, 2011. 4
- [6] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar. Fast and robust multiframe super resolution. *Image processing, IEEE Trans. on*, 13(10):1327–1344, 2004. 1
- [7] G. Freedman and R. Fattal. Image and video upscaling from local self-examples. *ACM Trans. on Graphics (TOG)*, 30(2):12, 2011. 2
- [8] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *Computer Graphics and Applications, IEEE*, 22(2):56–65, 2002. 2
- [9] D. Glasner, S. Bagon, and M. Irani. Super-resolution from a single image. In *ICCV*, pages 349–356. IEEE, 2009. 2
- [10] D. Krishnan and R. Fergus. Fast image deconvolution using hyper-laplacian priors. *Advances in Neural Information Processing Systems*, 22:1–9, 2009. 2
- [11] F. Malgouyres and F. Guichard. Edge direction preserving image zooming: a mathematical and numerical analysis. *SIAM Journal on Numerical Analysis*, 39(1):1–37, 2001. 2
- [12] Q. Shan, Z. Li, J. Jia, and C.-K. Tang. Fast image/video upsampling. In *ACM Trans. on Graphics (TOG)*, volume 27, page 153. ACM, 2008. 5
- [13] J. Sun, J. Sun, Z. Xu, and H.-Y. Shum. Gradient profile prior and its applications in image super-resolution and enhancement. *Image Processing, IEEE Trans. on*, 20(6):1529–1542, 2011. 2
- [14] Y.-W. Tai, S. Liu, M. S. Brown, and S. Lin. Super resolution using edge prior and single image detail synthesis. In *CVPR*, pages 2400–2407, 2010. 2
- [15] F. Viola, A. W. Fitzgibbon, and R. Cipolla. A unifying resolution-independent formulation for early vision. In *CVPR*, pages 494–501, 2012. 2
- [16] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492, 2010. 5, 7
- [17] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *Image Processing, IEEE Trans. on*, 19(11):2861–2873, 2010. 2, 5, 8
- [18] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006. 2
- [19] Z. Zhang, X. Liang, A. Ganesh, and Y. Ma. TILT: transform invariant low-rank textures. In *Computer Vision—ACCV 2010*, pages 314–328. Springer, 2011. 2, 3, 4




Figure 12. Results from super-resolving the images in Figure 8 using bicubic interpolation, total-variation regularization, sparse coding [17] and our proposed algorithm. The upsampling factor was 4 for sparse coding and 8 for the rest of the methods.



Figure 13. Super-resolved examples from Figure 12 obtained by TI-DTV regularization rectified by the transformation learnt by TILT from the low-resolution images.