

Image Denoising by Regularization on Characteristic Graphs

Thomas J. Asaki and Kevin R. Vixie

Department of Mathematics
103 Neill Hall
Washington State University
Pullman, WA 99164-3113

Matthew Sottile

Galois, Inc.
421 SW 6th Ave. Suite 300
Portland, OR 97204

Pavlo Cherepanov

Department of Mathematics and Statistics
MSC03 2150
University of New Mexico
Albuquerque, NM 87131-0001

Abstract

This paper introduces improvements to a now classical family of image denoising methods through rather minimal changes to the way derivatives are computed. In particular, we ask, and answer, the question “How much can we improve the common denoising methods by local, completely non-parametric modifications to image graphs?” We present the concept of *non-parametric* characteristic graph representations of images and detail two such graph constructions. Their use in image denoising is demonstrated within a regularization framework. The results are compared with those of more traditional approaches of Tikhonov, total variation and L^1TV regularization. We show that in some denoising scenarios our methods perform more favorably in preserving intensity levels and geometric details of object boundaries. They are particularly useful for denoising images with both smooth and discontinuous intensity variations, preserving detail to the pixel level.

Mathematics Subject Classification: 65D18, 68U10, 05C90.

Keywords: Image Restoration, Denoising, Graph Regularization.

1 Introduction

The methods of Mumford and Shah [16] and Rudin, Osher and Fatemi [19] are two of a small handful of image processing methods introduced about two decades ago which opened up the area of PDE and variational methods in image processing. Since then, an enormous amount of work has followed, inspired by their general approach.

A number of the most important approaches are special cases of

$$u^* = \arg \min_u (\|u - f\|_q^q + \alpha \|\nabla u\|_p^p), \quad (1)$$

where $f \in \mathbb{R}^n$ is a noisy image of n pixels, $1 \leq p, q < \infty$ (typically), and u and u^* are candidate and optimal denoised images, respectively. The data fidelity and regularization terms are weighted by the selection of a scalar α .

The popular H^1 realization of (1) ($p = 2, q = 2$) was first implemented in the work of Tikhonov (See [22]). The regularization enforces smoothness in u and is equivalent to a finite-time application of the heat equation to the noisy image. The net effect is to reduce all high-frequency content such as noise. One undesirable feature is that sharp intensity boundaries are smoothed.

The total variation (TV) realization (1) ($p = 1, q = 2$), or ROF model [19], gained rapid popularity from its ability to produce denoised images that retain sharp intensity boundaries. While successful for visual presentation, there are known side effects [20]. First, absolute intensity levels are altered. This is most obvious for images of piecewise constant intensity. Second, intensity regions of smooth variation become stair-stepped. Third, shapes in images are altered by elimination of high boundary curvature, reducing object perimeter.

While (1) has been highly successful and ubiquitously applied in image processing, the side effects arising from the regularization and choice of p are problematic for many applications. Several approaches have been taken in order to address these issues. Additional regularization based on object boundary lengths was proposed by Mumford and Shah [16]. This approach requires an additional functional weight parameter and assumes objects constrained by minimal perimeters.

Another approach proposed a variable p . See [4] and the developments in [13, 10, 5]. In this approach p is made a function of the gradient of the image. While this adds to the analytic complexity of the functional, the stair-stepping is reduced.

The L^1 TV approach ($p = 1, q = 1$), introduced in its continuous version by Chan and Esedoğlu in [8] after its initial discrete study by Alliney [3] and Nikolova [18], does a much better job of contrast preservation. This particular approach has generated continued interest; see for example [2, 15, 23].

Chartrand [9] has shown that solving the nonconvex optimization problem for $0 < p < 1$ yields improved results for shape and intensity preservation relative to TV. These results hold for images that can be represented sparsely in the gradient, that is, piecewise constant intensities.

Recently, *non-local means* and *diffusion geometric* methods [6, 12, 21] have generated significant improvements in denoising results by redefining what it means to

be a neighboring pixel. It is to the effect of image graph modifications that we now turn.

One way to view all the denoising methods is as operators on graphs where the values at vertices are pixel intensities and the edge weights are a function of some metric. Let G_0 denote the graph in which each pixel is connected to four neighbors (up, down, left, right). We will denote the fully connected graph, with $\frac{n(n-1)}{2}$ edges connecting the n pixels, by G . In the classical denoising case, the discrete version of (1) leads to p -Laplacian flows on G_0 . In the case of non-local means, edge weights might be taken to be $\frac{1}{k}$, where k is the number of nearest neighbors. The nearest neighbors are determined by considering distances in the space of pixel neighborhoods. This leads to a completely different graph which is non-local with respect to the natural metric suggested by the x-y image distance.

In this paper we take a different approach and ask how much improvement can be gained through very simple, even rather minimal perturbations to G_0 . We introduce a new method modifying the classical image graph G_0 through non-parametric decimation of the graph. This generates a subgraph $K \subset G_0$ on which we can use many of the classical and less classical denoising methods. The results show a marked improvement with respect to edge preservation, contrast preservation, and staircasing.

A critical point is that we are **not** examining what is the best we can do through a cleverly weighted subgraph of G , but rather what can be done using graphs K close to the usual graph G_0 . Accordingly, we compare our results to classical methods applied to images implicitly using G_0 , not the various subgraphs of G that the newer methods use.

Viewed another way, we simply remove edges from G_0 that connect *distant* pixels (but we do this in a non-parametric way) while the newer methods completely redefine what it means to be *close* (but in a parametric way). That our graph modifications are completely non-parametric, accomplished with no input other than the measured image, is an attractive feature that reduces the number of parameters that must be chosen to denoise an image.

Finally, we mention that because our goal is not to create the best denoising algorithm, but rather see what can be gained by minimal perturbation to rather classical methods, we do not consider various uses of sparse methods (but see for example [1]).

2 Approach

Consider the discretized image gradient of (1). Let $u_{i,j}$ be the image intensity at image pixel location (i, j) . A standard forward difference gradient approximation is

$$|\nabla u_{i,j}| \equiv \sqrt{(u_{i+1,j} - u_{i,j})^2 + (u_{i,j+1} - u_{i,j})^2}. \quad (2)$$

The first (second) difference term in (2) is the vertical (horizontal) gradient. We impose the Neumann boundary condition $\nabla u_{i,j} = 0$ if either i or j causes $u_{i,j}$ to fall

outside of the image.

This gradient can be represented as an operation on the graph $G_0 = (V_0, E_0)$ with vertex set $V_0 = \{v_{i,j}\}$ corresponding to all image pixels, and edge set $E_0 = \{e_k\}$ containing all pixel pairs that share a pixel edge (*not a graph edge*). We note that $v_{i,j}$ is a graph vertex, while $u_{i,j} = u(v_{i,j})$ is the image intensity at location $\{i, j\}$. This graph has been called the grid graph [7] and is recognized as that of a fully connected Von Neumann neighborhood (see [17]). Quite simply, G_0 is the fully-connected nearest (geographic) neighbor graph in the 2-d image space. Figure 1(b) illustrates G_0 for a small sample image. For image applications, we consider the weighted graph with edge weights given by $w_k = |u_{i_1, j_1} - u_{i_2, j_2}|$ for each edge $e_k = (v_{i_1, j_1}, v_{i_2, j_2})$, the absolute value of the image intensity difference associated with the vertex pair connected by the edge.

The denoising problem is now represented as

$$u^* = \arg \min_u (\|u - f\|_q^q + \alpha \|\nabla_K u\|_p^p). \quad (3)$$

where the gradient operator is applied to a graph K with weights W . The fundamental graph, $K = G_0$, is the one typically applied. We will consider subgraphs $K = \{V = V_0, E \subset E_0\} \subset G_0$ which we consider to be *characteristic* of images for denoising purposes. In the next section we will detail two such graphs.

The optimization problem described in (3) is nice: it is convex for $(1 \leq p, q \leq \infty)$. Though the regularization term suffers from a degeneracy yielding a null space with dimension equal to the number of connected components of the graph, the data fidelity term keeps the functional coercive.

In practice, the gradient computation is carried out over the fundamental graph G_0 with imposed zero edge weights that effectively define the subgraph of interest. More precisely,

$$\|\nabla_K u\|_p^p = \sum_{w_k \in W} |w_k|^p. \quad (4)$$

Note in particular that we are using the anisotropic version¹ of (3) here: this is the only version that makes sense on the graph – we do not have differences in the x and y directions at every node in the reduced graph K unless the graph is the complete graph G_0 .

We utilize the Lagged-Diffusivity method for solving (3) together with an application of the discrepancy principle for choosing the constant α [24]. We choose this approach for its general applicability to $p > 0, q > 0$, though we limit our presentation to three typical scenarios: Tikhonov regularization (H^1) [22] where $(p, q) = (2, 2)$; total variation regularization (TV) [19] where $(p, q) = (1, 2)$; and (L^1TV) [8] where $(p, q) = (1, 1)$. Lezoray, et al. [14] present a general method for solving (3) on graphs of arbitrary topology. They also present a variety of results for color image denoising using the fundamental graph G_0 . Their focus is on computational aspects of very general discrete graph-based denoising. Our main objective

¹In the case of $G = G_0$, the isotropic version is given by $\sum_i ((u_x^i)^2 + (u_y^i)^2)^{p/2}$ instead of the anisotropic analog given by $\sum_i (|u_x^i|^p + |u_y^i|^p)$

is the construction and application of characteristic graphs that are natural choices for image processing applications.

3 Characteristic Graphs

In this section we present two new characteristic image graphs, detail how they are constructed, and briefly discuss our expectations in denoising applications. An example of each graph is shown in Figure 1 as computed from a simple example image.




Figure 1: Left to right: A simple noisy test image; and the corresponding three characteristic graphs: the full graph or grid graph G_0 , the truncated Kruskal algorithm graph with next-nearest neighbor edge extension K_1 , and the vertex inclusion graph, K_2 .

In particular, our subgraphs are defined by features that are perturbations from the standard local graph structure. They have the following properties.

1. A graph should preserve spatial information. Distinct image regions should only be able to influence each other through neighbor pixel paths through the graph. By choosing graphs as subgraphs of G_0 this requirement is automatically met.
2. A graph should naturally detect image intensity content. In image regions of relatively small intensity variation, the graph should be highly connected. Conversely, graphs should have few or no edges corresponding to neighboring pixels with a large difference in intensity.
3. Graph construction should be non-parametric utilizing *only* the image intensity information without relying on user judgement or prior information.

These subgraph properties are intended to inspire *nonparametric* extensions to local denoising methods posed as in (1).

3.1 K1: Truncated-Kruskal Graph

We introduce a graph K_1 based on Kruskal's algorithm for constructing a minimum spanning tree of an arbitrary graph[11]. The tree construction is prematurely ter-

minated when all vertices are included and then additional graph edges are included to improve graph density.

1. Begin with an empty graph $V = \{\}$ and $E = \{\}$.
2. Add to E the edge of minimum weight in E_0 that does not create a cycle in (V, E) . Add the corresponding vertices to V that are not already in V .
3. Repeat step 2 until $V = V_0$.
4. Find the edge of minimum weight in E_0 that does not create a cycle in (V, E) . Set w_{cut} equal to the corresponding edge weight.
5. To each vertex of degree 1 add to E the associated edge from E_0 that (a) has smallest weight, (b) is not already in E , and (c) if the edge weight is less than w_{cut} .

The first four steps of the algorithm retain the time complexity of Kruskal's algorithm, $O(|E|\log|V|)$. The last step requires a $O(|E|\log|E|)$ sorting step to determine the minimum weight edge connected to each vertex. Due to the special structure of the grid graph derived from the image, we know that $|E|$ is bounded from above by $4|V|$, so this complexity is equivalent to that of Kruskal's algorithm.

The characteristic graph K1 will attempt to prevent graph connections across pixel neighbors of significantly different intensities. The density may be low relative to the grid graph G0, and the connectivity even in regions of like intensity may be circuitous or lacking. The third image in Figure 1 illustrates these properties. Note the somewhat circuitous graph paths one must often take to traverse from one pixel to its neighbor even if the intensities are similar. The best graph property is the absence of any graph edges that would connect object regions to background regions. It can be reasonably expected that this intensity discontinuity discrimination will be fairly robust to noise.

3.2 K2: Vertex Inclusion Graph

We also introduce the Vertex Inclusion Graph K2 in order to address the potential drawbacks of K1, low density and significant dis-connectivity. K2 is constructed by including all edges of weight less than or equal to a cutoff value non-parametrically determined from a vertex inclusion condition.

1. Begin with an empty graph $V = \{\}$ and $E = \{\}$.
2. Add the edge of smallest weight from E_0 to E that is not already in E , and the corresponding vertices from V_0 to V that are not already in V .
3. Repeat step 2 until $V = V_0$.
4. Add all edges from E_0 to E that are of equal weight to the largest edge weight in E .

Graph K2, like graph K1 can be disconnected across image intensity discontinuities. Also, we expect a significant number of graph edges in regions where the intensity variations are governed by noise. These properties are evident in the fourth image of Figure 1. Note both the absence of edges that cross the object-background boundary and the high graph density relative to K1. For very noisy images or images with salt and pepper noise, K2 can produce a nearly fully connected graph $K2 \sim G0$. In these cases, there is no expected advantage to using K2. Like the construction of K1, this approach has logarithmic time complexity of $O(|E| \log |E|)$ due to sorting the edges by edge weight.

3.3 Discussion on Parametrization

These graphs are considered to be *characteristic* representations of images because they satisfy, in our estimation, the criteria given at the beginning of this section. One might argue that these graph constructions contain hidden parameters. For example, why did we choose to add edges only to leaf vertices in forming K1? And, why did we specify only second nearest neighbors instead of, say, third nearest? Parameters such as these are of a different nature than more intrusive and application specific parameters such as edge weight thresholding, graph density thresholding, and local graph construction based on prior assumptions. We contend that our choices are related primarily to the desired structure of the graphs independent of the constituent intensity distributions of the images from which they are derived. Our graphs are intended for very general application because they capture spatial and intensity connectivity for a large class of images.

4 Results

We now present some specific image denoising results on several test images. The approach is to illustrate how the use of characteristic image graphs can provide improved intensity level and geometric content preservation relative to standard local methods.

Figure 2 shows denoising results for a simple image of piecewise constant intensities and several geometric features. The images at top are the clean image, a noisy image with additive Gaussian noise of $\sigma = 0.3$, and the color scale for all images in the figure. The clean image contains four overlapping objects at intensity values (1, 2, 3, 4) on a background of intensity zero. Nine denoised images are shown in the lower part of the figure, corresponding to the realizations of the three regularization scenarios (L^1TV , TV , H^1) applied using the three graphs ($G0$, $K1$, $K2$).

The piecewise constant intensities in this test image suggest the beneficial use of total variation based denoising. The standard result (TV - $G0$) attains relatively sharp boundaries and improved noise suppression over H^1 - $G0$. One residual effect is intensity stairstepping, seen as dark green areas bordering intensity transition regions. There is also some geometric information loss in regions of high curvature (small features), observed as boundary definition loss (a) between the dark green rectangle and the blue circle and (b) of the bottom of light blue triangle below

the magenta circle. The use of our characteristic graphs alleviates these problems. Both TV -K1 and TV -K2 nearly eliminate stairstepping and are much better at preserving the geometry of the original objects in the image. For this test image, L^1TV denoising produces very similar results to those of the corresponding TV results. Close examination reveals a greater degree of regional stairstepping with K1 over K2, as expected. It is also interesting to note how well H^1 methods perform using the graphs K1 and K2. Because these characteristic graphs were able to define lengthy connected stretches of object boundaries, a heat kernel smoothing approach is fairly successful in regional smoothing with significant sharp intensity boundary preservation.

Finally, we consider object intensity level recovery for the three TV based denoising methods. How accurate are the object intensities in denoised images? Table 1 lists the recovered pixel value mean and standard deviation over same the regions defined by the noiseless image. We find that intensity values are more accurately recovered using K1 or K2, especially for the highest and lowest values where G0 based methods are expected to fail. No method accurately recovers the intensity level of the smallest object feature at intensity 1. The characteristic graph based methods show noticeably smaller variances in intensity. This is due to their ability to significantly reduce stairstepping in intensity discontinuity regions.

G0	K1	K2
0.042(0.079)	0.003(0.028)	0.002(0.018)
0.669(0.264)	0.564(0.099)	0.515(0.116)
1.985(0.170)	2.036(0.055)	2.037(0.017)
2.957(0.110)	2.992(0.051)	3.007(0.073)
3.948(0.088)	3.994(0.023)	3.992(0.007)

Table 1: Intensity preservation comparison for TV denoising examples in Figure 2. The intensities in the noiseless figure are the integer values (0, 1, 2, 3, 4). This table gives the recovered pixel value mean and standard deviation over same the regions defined by the noiseless image.

Figure 3 shows denoising results for an image of piecewise smooth intensity variation with sharp boundaries. The images at top are the clean image, a noisy image with additive Gaussian noise of $\sigma = 0.16$, and the color scale for all images in the figure. The clean image contains a Gaussian bump with a sign change across a sinusoidal boundary. The intensity difference across the boundary approaches 2 near the image center and is much less than 2 near the left and right image boundary. Nine denoised images are shown in the lower part of the figure, corresponding to the realizations of the three regularization scenarios (L^1TV , TV , H^1) applied using the three graphs (G0, K1, K2).

This test image illustrates the dilemma of how to choose an appropriate regularization scenario when the image intensity is known to have both smooth variations and sharp discontinuities. An H^1 -G0 approach may denoise well and retain smooth

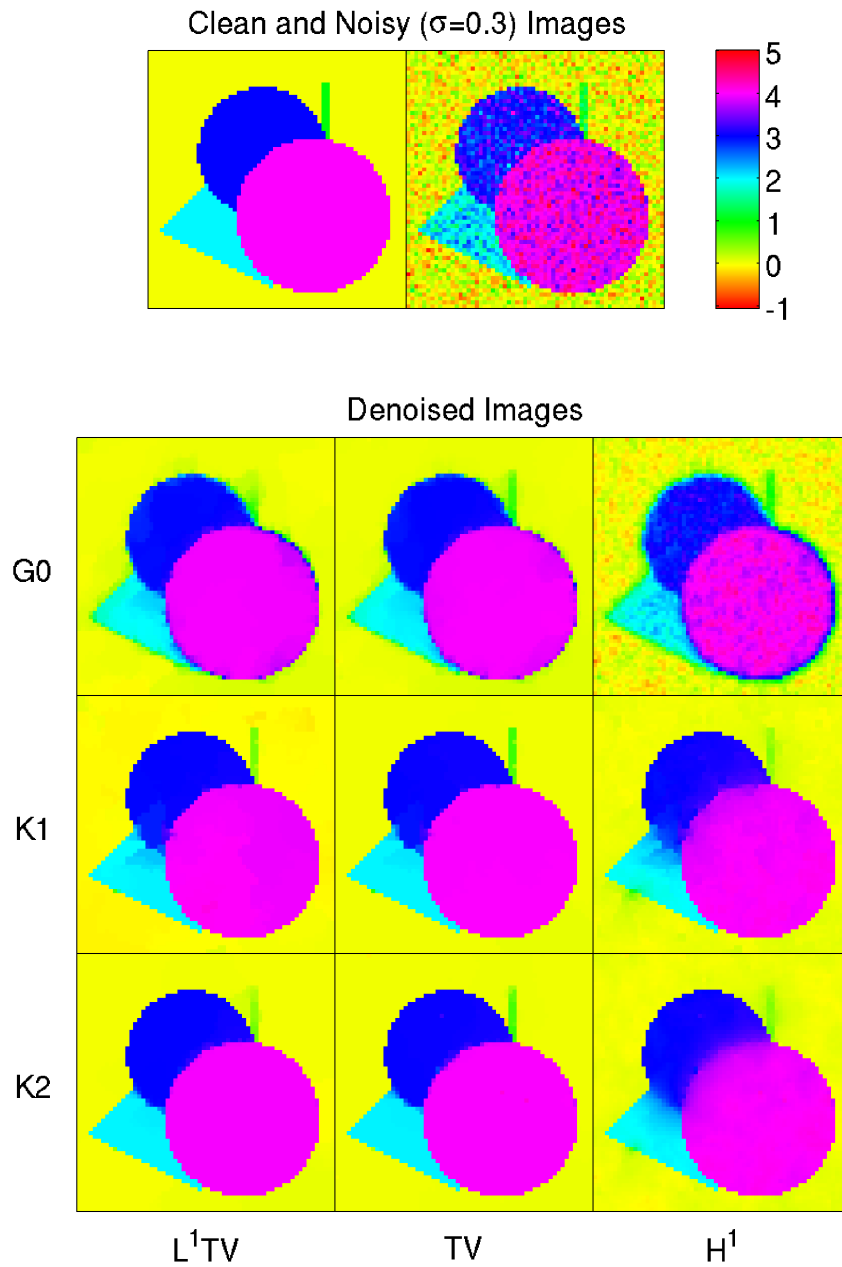


Figure 2: Denoising results for a simple image of piecewise constant intensities and several geometric features.

intensity variations, but at the cost of blurring discontinuities. A TV -G0 approach will perform much better at preserving the discontinuities, but will compromise the smoothly varying features, especially in the presence of noise. Our approach is to apply H^1 -K2 which does an excellent job of preserving both image features. The

denoised image is everywhere smoothly varying except across a well-defined sharp discontinuity boundary. The TV -K2 and L^1TV -K2 approaches have even better intensity discontinuity preservation but at the cost of regional intensity stairstepping. The K1 graph based results also show excellent intensity discontinuity preservation, though the image intensities show a large degree of regional stairstepping.

Figure 4 shows denoising results for an image of piecewise smooth intensity variation with sharp boundaries and high curvature geometric features. The images at top are the clean image, a noisy image with additive Gaussian noise of $\sigma = 0.12$, and the color scale for all images in the figure. The clean image contains a Gaussian bump with a unit jump within a star-shaped region. The star region has features of single pixel width. Nine denoised images are shown in the lower part of the figure, corresponding to the realizations of the three regularization scenarios (L^1TV , TV , H^1) applied using the three graphs (G0, K1, K2).

This test image is similar to the previous image in that the intensities are smoothly varying except across some discontinuous boundary features. However, the object now has very narrow features. We would like to preserve (a) smooth intensity variation, (b) sharp intensity discontinuities, and (c) geometric features at all length scales. The standard H^1 -G0 approach succeeds only at the first. The TV -G0 approach succeeds only at the second. The L^1TV -G0 approach also preserves sharp intensity discontinuities, but dramatically fails to preserve the geometry and small length scales. Our H^1 -K2 approach succeeds on all points. The other denoising results are shown for comparison, though it is notable that the K1 and K2 graph methods are very good at preserving geometric information in this particular low noise example.

Figure 5 shows denoising results for the same images as in Figure 4 but with a larger noise amplitude ($\sigma = 0.4$). The images at top are the clean image, a noisy image, and the color scale for all images in the figure. Nine denoised images are shown in the lower part of the figure, corresponding to the realizations of the three regularization scenarios (L^1TV , TV , H^1) applied using the three graphs (G0, K1, K2).

Under these higher noise conditions, the standard G0 based methods perform similarly to the previous case (Figure 4) but with greater object distortion, both in loss of small length scale features and object boundary preservation. The K2 based results show somewhat improved object boundary preservation but are qualitatively similar to the G0 results. This is because the K2 graph becomes nearly complete in the presence of noise of standard deviation comparable to half the size of the actual intensity discontinuity. The K1 based methods do a better job of object boundary preservation because the graph construction is more robust to noise. The TV -K1 and L^1TV -K1 results show the best object/background discrimination including some of the small length scale features.

Figure 6 shows denoising results for a silhouette image with objects containing a variety of length scales. The images at top are the clean image, a noisy image with additive Gaussian noise of $\sigma = 0.15$, and the color scale for all images in the figure. The clean image contains a binary image of berries and branches with important length scales from one to 20 pixels. Nine denoised images are shown in the lower part

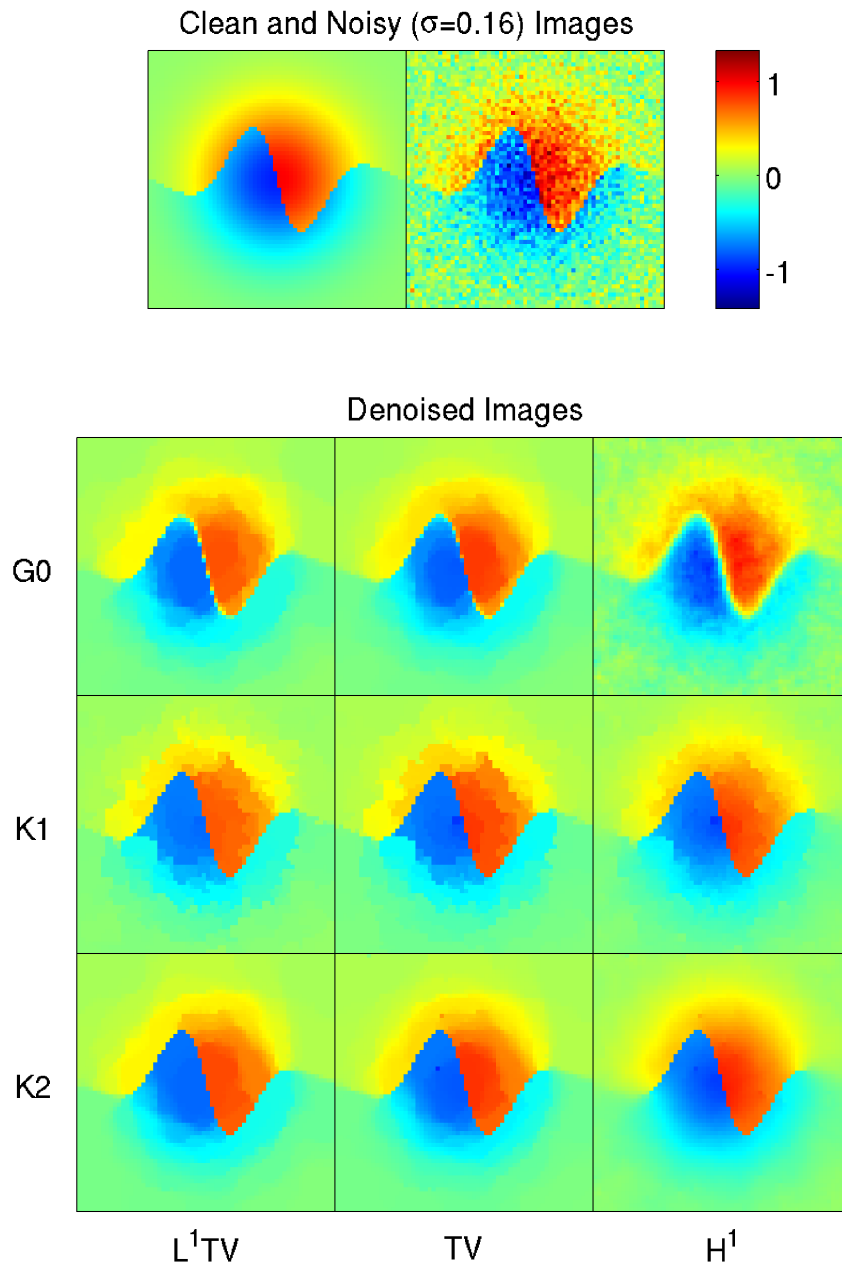


Figure 3: Denoising results for an image of piecewise smooth intensity variation with sharp boundaries.

of the figure, corresponding to the realizations of the three regularization scenarios (L^1TV , TV , H^1) applied using the three graphs (G0, K1, K2).

The TV based results illustrate the improved sharpness of the boundary representation with K1 and K2 over the typical G0. Once again this is a consequence of




Figure 4: Denoising results for an image of piecewise smooth intensity variation with sharp boundaries and high curvature geometric features – low noise realization.

K1 and K2 having few graph connections across the object/background boundary. The contrast retention is better as well, the differences between the pixel mean values between object and background for each case are: G0, 0.935; K1, 0.979; K2, 0.981. The feature retention with L^1TV is better using K1 or K2 relative to G0; note the

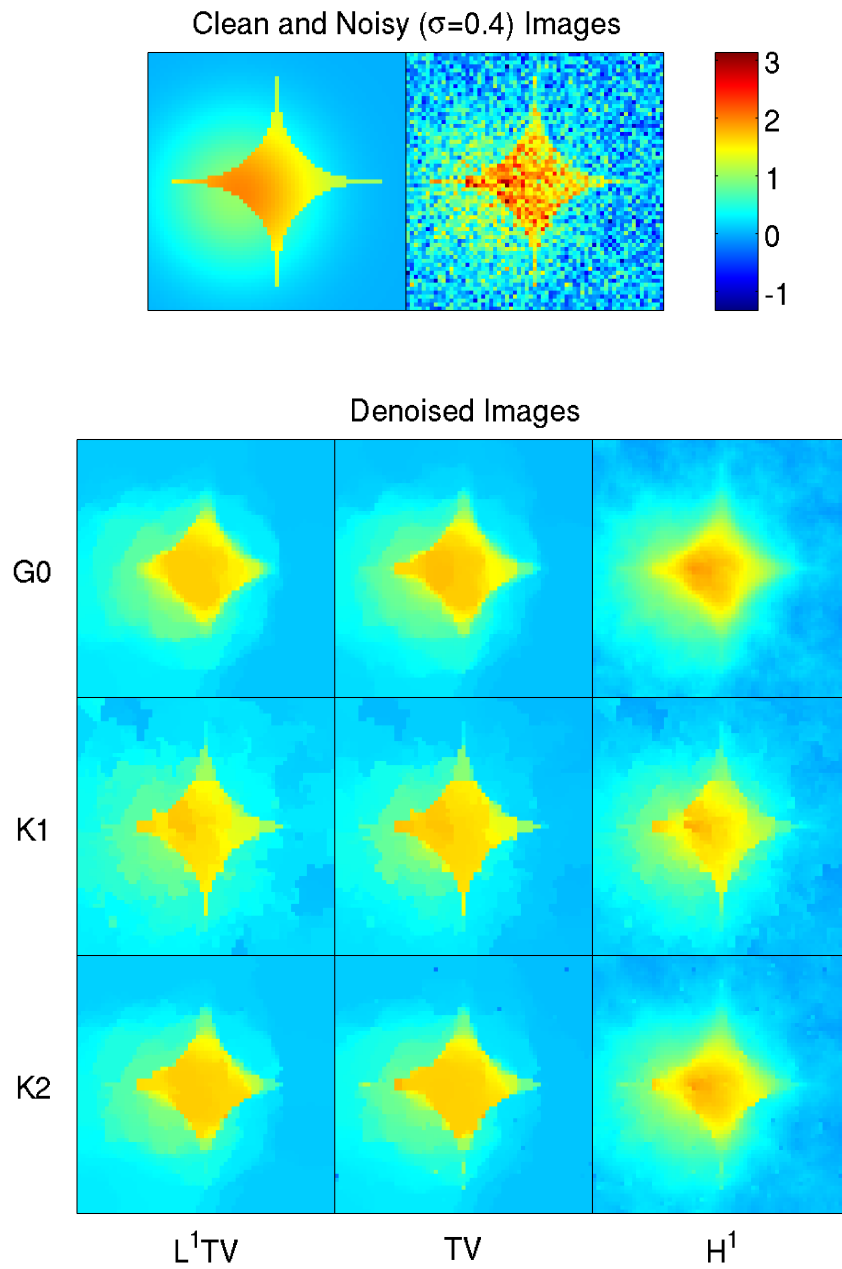


Figure 5: Denoising results for an image of piecewise smooth intensity variation with sharp boundaries and high curvature geometric features– high noise realization.

disappearance of the narrowest braches in L^1TV -G0. It is worth noting that both K1 and K2 graphs are not completely disconnected across all object/background boundaries. This is most evident in examining H^1 -K1 and H^1 -K2, in which the areas of boundary crossings are seen by the yellowish diffusion “clouds”.

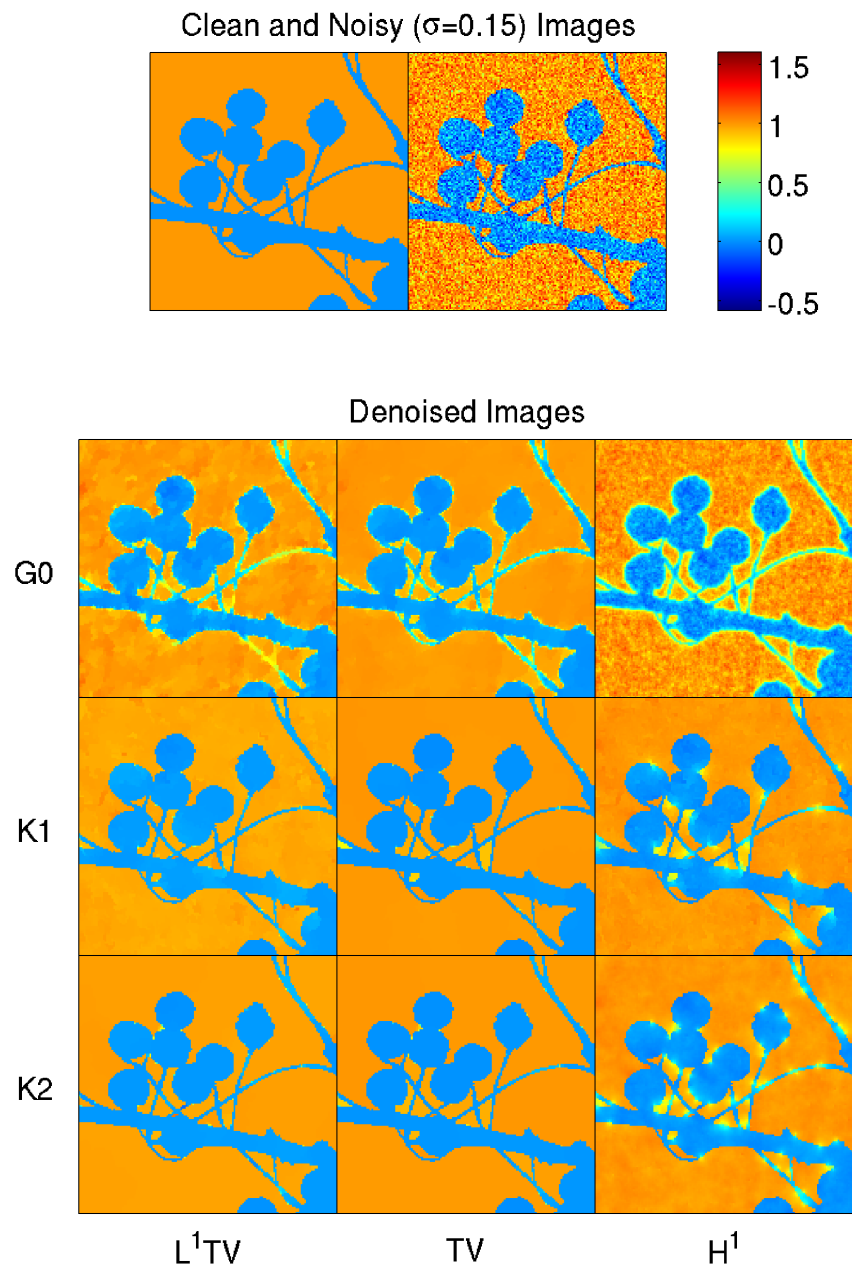


Figure 6: Denoising results for a silhouette image with objects containing a variety of length scales – low noise realization.

Figure 7 shows denoising results for the same images as in Figure 6 but with a larger noise amplitude ($\sigma = 0.40$). The images at top are the clean image, a noisy image, and the color scale for all images in the figure. Nine denoised images are shown in the lower part of the figure, corresponding to the realizations of the three

regularization scenarios (L^1TV , TV , H^1) applied using the three graphs (G0, K1, K2).

In this higher noise example we observe that the K2 based results are similar to the G0 based results, again due to the high connectedness of K2 in high noise situations. In these cases, there is good object feature recovery for the largest length scales (berries and large branch) and noticeably worse recovery for the smallest length scales (small branches). The small length scale loss is most pronounced in the L^1TV cases. However, the K1 based results, while similar in other respects, show better recovery of the geometry of small length scale features relative to the G0 and K2 counterparts. This is a consequence of the K1 graph construction being more robust to image noise.

5 Discussion

The characteristic graphs K1 and K2 are especially useful for denoising images of piecewise smooth content. They help to preserve intensity levels because intensity jumps are not penalized if they are large relative to a characteristic weight determined by an appropriate construction algorithm. They also help to preserve object boundary shape details for much the same reason. Intensity jump boundaries have zero penalty regardless of length or curvature. Object details can be preserved down to the pixel level. Even in the non-ideal case, where noise has allowed graph construction algorithms to build links across intensity discontinuities, these links are sparse relative to a full graph implementation and still serve to reduce distortion effects.

Our characteristic graphs will fail to produce significantly improved results in some cases. If images are very noisy, are corrupted by salt and pepper noise, or contain intensity isolated pixels, then we expect the characteristic K2 graph to be nearly as dense as the full G0 graph. This simply means that the image information needed to construct a good subgraph, $K2 \subset G0$, is lacking. The K1 graph construction is more robust to image noise, but no graph construction method presented here is expected to perform well when noise amplitude is comparable to discontinuous intensity jumps within the true image. We note that our results are not expected to be inferior to methods that employ G0.

The use of characteristic image subgraphs is general to PDE-based methods for image processing. We have defined here example graphs and demonstrated their potential using some simple denoising examples. Clear extensions are applications to segmentation and texture extraction.

We have focused exclusively on non-parametric modifications to (1). Other important graphs are certainly possible that make use of (problem-dependent) parameters. Some examples include graphs that achieve a certain density or connectivity, or user-defined graph construction cutoff values. We also note that there are certain possible parametric modifications to the denoising procedure. For example, one is tempted to employ a variable α procedure in which the characteristic graph is updated as α is incrementally raised from some low value to a final value. This may help to reduce graph connectivity across boundaries significantly obscured by noise.




Figure 7: Denoising results for a silhouette image with objects containing a variety of length scales – high noise realization.

6 Conclusion

We have defined two characteristic image graphs, the truncated Kruskal graph K1 and the vertex inclusion graph K2 as nonparametric extensions to local PDE im-

age denoising methods. We have demonstrated their use by application to several simple test images. Comparisons were made with standard TV , L^1TV , and H^1 denoising methods. Our characteristic graph based denoising methods demonstrate (a) improved contrast preservation, (b) improved geometric feature recovery even down to the pixel level, and (c) simultaneous recovery of smooth intensity variation and discontinuous intensity jumps. The characteristic graph constructions are stable and predictable with respect to noise type and amplitude. The K2 graph closely approximates the full mesh graph G_0 in cases of very high noise or corruption by salt and pepper noise. The K1 graph remains relatively sparsely connected in all noise situations but may fail to be meaningful when noise amplitude is comparable to image feature intensity differences.

ACKNOWLEDGEMENTS. This research was performed, in part, under the auspices of the Department of Energy under contract W-7405-ENG-36. It was initially funded at Los Alamos National Laboratory through the 2007 Data Driven Modeling and Analysis summer school jointly funded by the T-07 summer student program and the New Mexico Consortium.

References

- [1] M. AHARON, M. ELAD, AND A.M. BRUCKSTEIN, *The K-SVD: An Algorithm for Designing of Overcomplete Dictionaries for Sparse Representation*, the IEEE Trans. On Signal Processing, Vol. 54, no. 11, pp. 4311-4322, November 2006.
- [2] W. K. ALLARD, *Total variation regularization for image denoising; I. Geometric Theory.*, To appear in SIAM Journal on Mathematical Analysis, (2007).
- [3] S. ALLINEY, *A property of the minimum vectors of a regularizing functional defined by means of the absolute norm*, IEEE Trans. Signal Process., 45 (1997), pp. 913-917.
- [4] P. BLOMGREN, T. F. CHAN, P. MULET, AND C. WONG, *Total variation image restoration: Numerical methods and extensions*, in Proceedings of the 1997 IEEE International Conference on Image Processing, vol. III, 1997, pp. 384-387. CAM report 97-50.
- [5] E. M. BOLLT, R. CHARTRAND, S. ESEDOĞLU, P. SCHULTZ, AND K. R. VIXIE, *Existence and non-uniqueness for $\int |\nabla u|^{p(|\nabla u|)}$ regularized image functionals*, Advances in Computational Mathematics, Special issue, "Mathematical Methods for Image Processing" (2008).
- [6] A. BUADES, B. COLL, AND J. M. MOREL, *A Review of Image Denoising Algorithms, with a New One*, Multiscale Modeling and Simulation, vol. 4, no. 2, pp. 490-530, July 2005.
- [7] T. CHAN, S. OSHER, AND J. SHEN, *The digital TV filter and nonlinear denoising*, IEEE Transactions on Image Processing, 10 (2001).

- [8] T. F. CHAN AND S. ESEDOĞLU, *Aspects of total variation regularized L^1 function approximation*, SIAM J. Appl. Math., 65 (2005), pp. 1817–1837.
- [9] R. CHARTRAND, *Nonconvex regularization for shape preservation*, in IEEE International Conference on Image Processing (ICIP), 2007.
- [10] Y. CHEN, S. LEVINE, AND M. RAO, *Variable exponent, linear growth functionals in image restoration*, SIAM Journal of Applied Mathematics, 66 (2006), pp. 1383–1406.
- [11] T. H. CORMEN, C. E. LEISERSON, R. L. RIVEST, AND C. STEIN, *Introduction to Algorithms*, The MIT Press, second ed., 2002.
- [12] K. DABOV, A. FOI, V. KATKOVNIK, AND K. EGIAZARIAN, *Image denoising by sparse 3D transform-domain collaborative filtering*, IEEE Transactions on Image Processing, vol. 16, no. 8, August 2007.
- [13] S. E. LEVINE, *An adaptive variational model for image decomposition*, in Energy Minimization Methods in Computer Vision and Pattern Recognition, no. 3757 in Lecture Notes in Computer Science, Springer, 2005, pp. 382–397.
- [14] O. LEZORAY, A. ELMOATAZ, AND S. BOUGLEUX, *Graph regularization for color and image processing*, Computer Vision and Image Understanding, 107 (2007).
- [15] S. P. MORGAN AND K. R. VIXIE, *L^1 TV computes the flat norm for boundaries*, Abstract and Applied Analysis, 2007 (2007), pp. Article ID 45153, 14 pages. doi:10.1155/2007/45153.
- [16] D. MUMFORD AND J. SHAH, *Optimal approximations by piecewise smooth functions and associated variational problems*, Communications on Pure and Applied Mathematics, 42 (1989).
- [17] J. V. NEUMANN AND A. W. BURKS, *Theory of self-reproducing automata*, University of Illinois Press, 1966.
- [18] M. NIKOLOVA, *Minimizers of cost-functions involving nonsmooth data-fidelity terms*, SIAM J. Numer. Anal., 40 (2003), pp. 965–994.
- [19] L. RUDIN, S. OSHER, AND D. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, 60 (1992).
- [20] D. STRONG AND T. CHAN, *Edge-preserving and scale-dependent properties of total variation regularization*, Inverse Problems, 19 (2003).
- [21] A.D. SZLAM, M. MAGGIONI, R.R. COIFMAN *Regularization on graphs with function adapted diffusion processes*, Journal of Machine Learning Research, 9: 1711–1739, 2008.
- [22] A. N. TIKHONOV AND V. Y. ARSEININ, *Solutions of ill-posed problems*, Winston, 1977.

- [23] K. R. VIXIE, *Some properties of minimizers for the Chan-Esedoğlu L^1 TV functional*, arXiv.org, (2007).
- [24] C. VOGEL, *Computational methods for inverse problems*, SIAM, 2002.

Received: Month xx, 200x