

Improving Density Estimation by Incorporating Spatial Information

Laura M. Smith Matthew S. Keegan Todd Wittman George O. Mohler
lsmith@math.ucla.edu mkeegan@math.ucla.edu wittman@math.ucla.edu gmohler@math.ucla.edu
Andrea L. Bertozzi
bertozzi@math.ucla.edu

Department of Mathematics, University of California, Los Angeles

November 30, 2009

Abstract

Given discrete event data, we wish to produce a probability density that can model the relative probability of events occurring in a spatial region. Common methods of density estimation, such as Kernel Density Estimation, do not incorporate geographical information. Using these methods could result in non-negligible portions of the support of the density in unrealistic geographic locations. For example, crime density estimation models that do not take geographic information into account may predict events in unlikely places such as oceans, mountains, etc. We propose a set of Maximum Penalized Likelihood Estimation methods based on Total Variation and H_1 Sobolev norm regularizers in conjunction with a priori high resolution spatial data to obtain more geographically accurate density estimates. We apply this method to a residential burglary data set of the San Fernando Valley using geographic features obtained from satellite images of the region and housing density information.

1 Introduction

High resolution and hyperspectral satellite images, city and county boundary maps, census data, and other types of geographical data provide much information about a given region. It is desirable to integrate this knowledge into models defining geographically dependent data. Given spatial event data, we will be constructing a probability density that estimates the probability that an event will occur in a region. Often, it is unreasonable for events to occur in certain regions, and we would like our model to reflect this restriction. For example, residential burglaries and other types of crimes are unlikely to occur in oceans, mountains, and other regions. Such areas can be determined using aerial images or other external spatial data, and we denote these improbable locations as the invalid region. Ideally, the support of our density should be contained in the valid region.

Geographic profiling, a related topic, is a technique used to create a probability density from a set of crimes by a single individual to predict where the individual is likely to live or work [15]. Some law enforcement agencies currently use software that makes predictions in unrealistic geographic locations. Methods that incorporate geographic information have recently been proposed and is an active area of research [11, 13].

A common method for creating a probability density is to use Kernel Density Estimation [18, 19], which approximates the true density by a sum of kernel functions. A popular choice for the kernel is the Gaussian distribution which is smooth, spatially-symmetric, and has non-compact support. Other probability density estimation methods include the taut string, logspline, and the Total Variation Maximum Penalized Likelihood Estimation models [3, 8, 17, 7, 10]. However, none of these methods utilize information from external spatial data. Consequently, the density estimate typically has some nonzero probability of events occurring in the invalid region. Figure 1 demonstrates these problems with the current methods and how the methods we will propose in this paper resolve them. Located in the middle of the image are two disks where events cannot occur, depicted in figure 1(a). We selected randomly from the region outside the disks using a predefined probabilistic density that is provided in figure 1(b). The 4,000 events chosen are shown in figure 1(c). With a variance of $\sigma = 2.5$, we see in figure 1(d) that the Kernel Density Estimate predicts that events may occur in our invalid region.

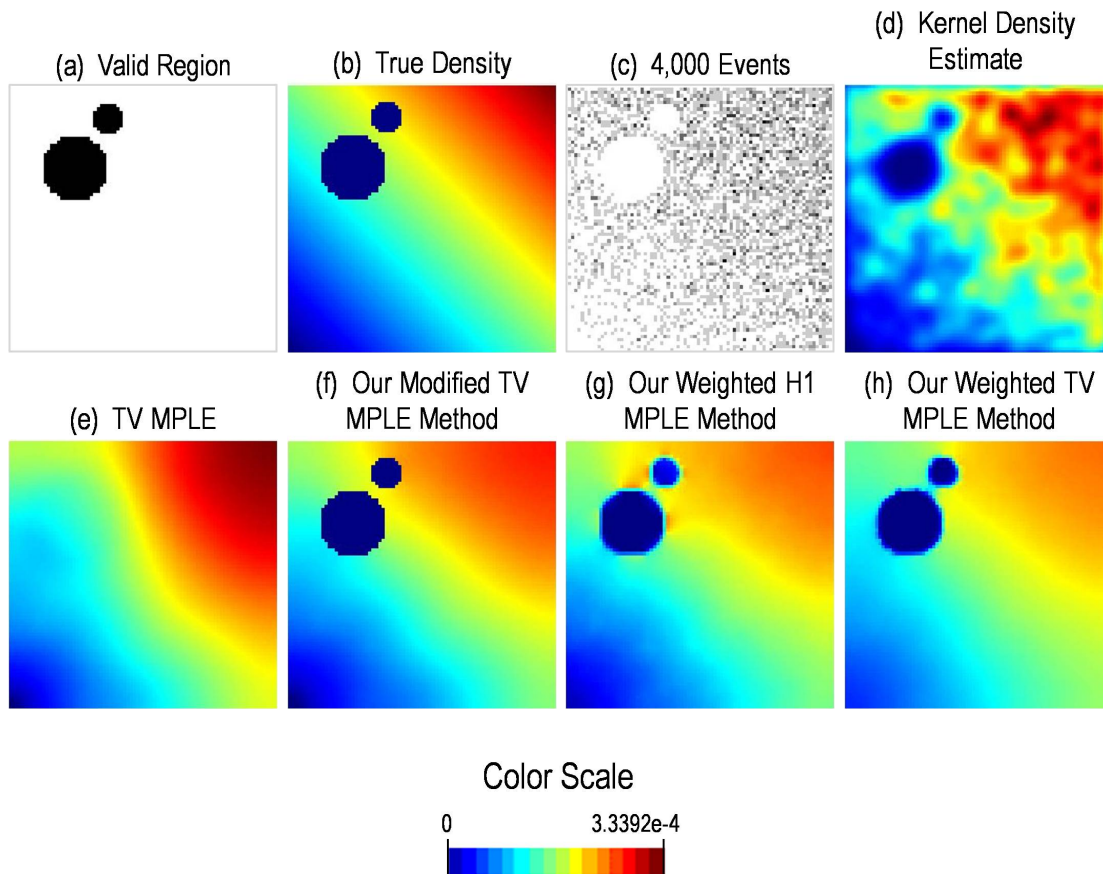


Figure 1: Simple Example with Comparison of Methods

In this article we propose a novel set of models that restrict the support of the density estimate to the valid region and ensure realistic behavior. The models use Maximum Penalized Likelihood Estimation [4, 6], which is a variational approach. The density estimate is calculated as the minimizer of some predefined energy functional. The novelty of our approach is in the way we define the energy functional with explicit dependence on the valid region such that the density estimate obeys our assumptions of its support. The results from our methods for this simple example are illustrated in figures 1(f), 1(g), and 1(h).

The article is structured in the following way. In section 2 Maximum Penalized Likelihood Methods are introduced. In sections 3 and 4 we present our set of models which we name the *Modified* Total Variation MPLE model and the *Weighted* H_1 Sobolev MPLE model, respectively. In section 5 we discuss the implementation and numerical schemes that we use to solve for the solutions of the models. We provide examples for validation of the models and an example with actual residential burglary data in section 6. In this section, we also compare our results to the Kernel Density Estimation model and other Total Variation MPLE methods. Finally, we discuss our conclusions and future work in section 7.

2 Maximum Penalized Likelihood Estimation

Assuming that $u(\mathbf{x})$ is the desired probability density for $\mathbf{x} \in \mathbb{R}^2$, and the known location of events occur at $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$, then Maximum Penalized Likelihood Estimation (MPLE) models are given by

$$\hat{u}(\mathbf{x}) = \underset{\int_{\Omega} u dx=1, 0 \leq u}{\operatorname{argmax}} \left\{ \sum_{i=1}^n \log(u(\mathbf{x}_i)) - \mu_0 P(u) \right\}.$$

Here, $P(u)$ is a penalty functional, which is generally designed to produce a smooth density map. This maximization problem can be written equivalently as

$$\hat{u}(\mathbf{x}) = \underset{\int_{\Omega} u dx=1, 0 \leq u}{\operatorname{argmin}} \left\{ P(u) - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right\}.$$

Here, the parameter μ determines how strongly weighted the maximum likelihood term is compared to the penalty functional.

A range of penalty functionals have been proposed, including $P(u) = \int_{\Omega} |\nabla \sqrt{u}|^2 dx$ [6, 4] and $P(u) = \int_{\Omega} |\nabla^3(\log(u))|^2 dx$ [18, 4]. More recently, variants of the Total Variation (TV) functional [16], $P(u) = \int_{\Omega} |\nabla u| dx$, have been proposed for MPLE [7, 17, 10]. These methods do not explicitly incorporate the information that can be obtained from the external spatial data, although some note the need to allow for various domains. Even though the TV functional will maintain sharp gradients, the boundaries of the constant regions do not necessarily agree with the boundaries within the image. This method also performs poorly when the data is too sparse, as the density is smoothed to have equal probability almost everywhere. Figure 1(e) demonstrates this, in addition to how this method predicts events in the invalid region with non-negligible estimates.

The methods we propose use a penalty functional that depends on the valid region determined from the geographical images or other external spatial data. Figure 1 demonstrates how these models will improve on the current methods.

3 The Modified Total Variation MPLE Model

The first model we propose is an extension of the Maximum Penalized Likelihood Estimation method given by Mohler, et al. [10],

$$\hat{u}(\mathbf{x}) = \operatorname{argmin}_{\int_{\Omega} u d\mathbf{x}=1, 0 \leq u} \left\{ \int_{\Omega} |\nabla u| d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right\}.$$

Once we have determined a valid region, we wish to align the level curves of the density function u with the boundary of the valid region. The Total Variation functional is well known to allow discontinuities in its minimizing solution [16]. By aligning the level curves of the density function with the boundary, we encourage a discontinuity to occur there to keep the density from smoothing into the invalid region.

Since $\frac{\nabla u}{|\nabla u|}$ gives the unit normal vectors to the level curves of u , we would like

$$\frac{\nabla(1_D)}{|\nabla(1_D)|} = \frac{\nabla u}{|\nabla u|},$$

where (1_D) is the characteristic function of the valid region D . The region D is obtained from external spatial data, such as aerial images. To avoid division by zero, we use $\theta := \frac{\nabla(1_D)}{|\nabla(1_D)|_\epsilon}$, where $|\nabla v|_\epsilon = \sqrt{v_x^2 + v_y^2 + \epsilon^2}$. To align the density function and the boundary one would want to minimize $|\nabla u| - \theta \cdot \nabla u$. Integrating this and applying integration by parts, we obtain the term $\int_{\Omega} |\nabla u| + u \nabla \cdot \theta d\mathbf{x}$. We propose the following *Modified* Total Variation penalty functional, where we adopt the more general form of the above functional

$$\hat{u}(\mathbf{x}) = \operatorname{argmin}_{\int_{\Omega} u d\mathbf{x}=1, 0 \leq u} \left\{ \int_{\Omega} |\nabla u| d\mathbf{x} + \lambda \int_{\Omega} u \nabla \cdot \theta d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right\}.$$

The parameter λ allows us to vary the strength of the alignment term. Two pan-sharpening methods, P+XS and Variational Wavelet Pan-sharpening [9, 2], both include a similar term in their energy functional to align the level curves of the optimal image with the level curves of the high resolution pan-chromatic image.

4 The Weighted H_1 Sobolev MPLE Model

A Maximum Penalized Likelihood Estimation method with penalty functional $\int_{\Omega} \frac{1}{2} |\nabla u|^2 d\mathbf{x}$, the H_1 Sobolev norm, gives results equivalent to those obtained using Kernel Density Estimation [4].

We enforce the H_1 regularizer term away from the boundary of the invalid region. This results in the model

$$\hat{u}(\mathbf{x}) = \operatorname{argmin}_{\int_{\Omega} u d\mathbf{x}=1, 0 \leq u} \left\{ \frac{1}{2} \int_{\Omega \setminus \partial D} |\nabla u|^2 d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right\}.$$

This new term is essentially the smoothness term from the Mumford-Shah model [12]. We approximate the H_1 term by introducing the Ambrosio-Tortorelli approximating function $z_{\epsilon}(\mathbf{x})$ [1], where $z_{\epsilon} \rightarrow (1 - \delta(\partial D))$ in the sense of distributions. More precisely, we use a continuous function which has the property

$$z_{\epsilon}(\mathbf{x}) = \begin{cases} 1 & \text{if } d(\mathbf{x}, \partial D) > \epsilon, \\ 0 & \text{if } \mathbf{x} \in \partial D. \end{cases}$$

Thus, the minimization problem becomes

$$\hat{u}(\mathbf{x}) = \operatorname{argmin}_{\int_{\Omega} u d\mathbf{x}=1, 0 \leq u} \left\{ \frac{1}{2} \int_{\Omega} z_{\epsilon}^2 |\nabla u|^2 d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right\}.$$

The weighting away from the edges is used to control the diffusion into the invalid region. This method of weighting away from the edges can also be used with the Total Variation functional in our first model, and we will refer to this as our *Weighted* TV MPLE model.

5 Implementation

5.1 The Constraints

In the implementation for the Modified Total Variation MPLE method and Weighted H_1 MPLE method, we must enforce the constraints $0 \leq u(\mathbf{x})$ and $\int_{\Omega} u(\mathbf{x}) d\mathbf{x} = 1$ to ensure that $u(\mathbf{x})$ is a probability density estimate. We will solve these minimization problems by obtaining their Euler-Lagrange equations, which are elliptic partial differential equations. These types of equations have maximum principles, so if our initial approximation has values greater than or equal to zero, then our final solution will as well. Thus, we can ignore the $u \geq 0$ constraint since it will be satisfied implicitly by the models.

We enforce the second constraint by first adding it to the energy functional as an L_2 penalty term. For the H_1 method, this change results in the new minimization problem

$$\hat{u}_H(\mathbf{x}) = \operatorname{argmin}_u \left\{ \frac{1}{2} \int_{\Omega} z_{\epsilon}^2 |\nabla u|^2 d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} - 1 \right)^2 \right\},$$

where we have denoted $\hat{u}_H(\mathbf{x})$ as the solution of the H_1 model. The constraint is then enforced by applying Bregman iteration [14]. Using this method, we formulate our problem as

$$(u_H, b_H) = \operatorname{argmin}_{u, b} \left\{ \frac{1}{2} \int_{\Omega} z_{\epsilon}^2 |\nabla u|^2 d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b - 1 \right)^2 \right\}$$

where b is introduced as the Bregman variable of the sum to unity constraint. We solve this problem using alternating minimization, updating the u and the b iterates as

$$(H1) \quad \begin{cases} u^{(k+1)} &= \operatorname{argmin}_u \left\{ \frac{1}{2} \int_{\Omega} z_{\epsilon}^2 |\nabla u|^2 d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b^{(k)} - 1 \right)^2 \right\} \\ b^{(k+1)} &= b^{(k)} + \int_{\Omega} u^{(k+1)} d\mathbf{x} - 1, \end{cases}$$

with $b^{(0)} = 0$. Similarly for the modified TV method, we solve the alternating minimization problem

$$(TV) \quad \begin{cases} u^{(k+1)} &= \operatorname{argmin}_u \left\{ \int_{\Omega} |\nabla u| d\mathbf{x} + \lambda \int_{\Omega} u \nabla \cdot \theta d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) \right. \\ &\quad \left. + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b^{(k)} - 1 \right)^2 \right\} \\ b^{(k+1)} &= b^{(k)} + \int_{\Omega} u^{(k+1)} d\mathbf{x} - 1 \end{cases}$$

with $b^{(0)} = 0$.

5.2 Weighted H_1 MPLE Implementation

For the Weighted H_1 MPLE model, the Euler-Lagrange equation for the u minimization is given by

$$(H1) \quad -\nabla (z_{\epsilon}^2 \nabla u) - \mu \frac{w(\mathbf{x})}{u(\mathbf{x})} + \gamma \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b^{(k)} - 1 \right) = 0,$$

where $w(\mathbf{x})$ is the given number of sampled events that occurred at \mathbf{x} . We solve this using a Gauss-Seidel method with central differences for the ∇z^2 and ∇u . Once we have discretized the partial differential equation, solving this equation simplifies to solving the quadratic

$$(4z^2 + \gamma) u_{i,j}^2 - \alpha_{i,j} u_{i,j} - \mu w_{i,j} = 0$$

for the positive root, where

$$\begin{aligned} \alpha_{i,j} &= z_{i,j}^2 (u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}) \\ &+ \left(\frac{z_{i+1,j}^2 - z_{i-1,j}^2}{2} \right) \left(\frac{u_{i+1,j} - u_{i-1,j}}{2} \right) \\ &+ \left(\frac{z_{i,j+1}^2 - z_{i,j-1}^2}{2} \right) \left(\frac{u_{i,j+1} - u_{i,j-1}}{2} \right) \\ &+ \gamma \left(1 - b^{(k)} - \sum_{(i',j') \neq (i,j)} u_{i',j'} \right). \end{aligned}$$

We chose our parameters μ and γ so that the Gauss-Seidel solver will converge. In particular, we have $\mu = O((NM)^{-2})$ and $\gamma = O(\mu(NM))$, where the image is $N \times M$.

5.3 Modified TV MPLE Implementation

There are many approaches for handling the minimization of the Total Variation penalty functional. A fast and simple method for doing this is to use the Split Bregman technique (see [5],[10] for an

in depth discussion, see also [20]). In this approach, we substitute the variable \mathbf{d} for ∇u in the TV norm and then enforce the equality $\mathbf{d} = \nabla u$ using Bregman iteration. To apply Bregman iteration, we introduce the variable \mathbf{g} as the Bregman vector of the $\mathbf{d} = \nabla u$ constraint. This results in a minimization problem in which we minimize with respect to both \mathbf{d} and u .

Beginning the iteration with $\mathbf{g}^{(0)} = 0$, the minimization is written as

$$\begin{aligned} (u^{(k+1)}, \mathbf{d}^{(k+1)}) &= \operatorname{argmin}_{u, \mathbf{d}} \left\{ \|\mathbf{d}\|_1 + \lambda \int_{\Omega} u \nabla \cdot \theta d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b^{(k)} - 1 \right)^2 \right. \\ &\quad \left. + \frac{\alpha}{2} \|\mathbf{d} - \nabla u - \mathbf{g}^{(k)}\|_2^2 \right\} \\ \mathbf{g}^{(k)} &= \mathbf{g}^{(k-1)} + \nabla u^{(k)} - \mathbf{d}^{(k)}. \end{aligned}$$

Alternating the minimization of $u^{(k+1)}$ and $\mathbf{d}^{(k+1)}$, we obtain our final formulation for the TV model as

$$(TV) \begin{cases} u^{(k+1)} &= \operatorname{argmin}_u \left\{ \lambda \int_{\Omega} u \nabla \cdot \theta d\mathbf{x} - \mu \sum_{i=1}^n \log(u(\mathbf{x}_i)) + \frac{\gamma}{2} \left(\int_{\Omega} u(\mathbf{x}) d\mathbf{x} + b^{(k)} - 1 \right)^2 \right. \\ &\quad \left. + \frac{\alpha}{2} \|\mathbf{d}^{(k)} - \nabla u - \mathbf{g}^{(k)}\|_2^2 \right\} \\ d_j^{(k+1)} &= \operatorname{shrink} \left((\nabla u^{(k+1)})_j - d_j^{(k)}, \frac{1}{\alpha} \right) \\ \mathbf{g}^{(k+1)} &= \mathbf{g}^{(k)} + \nabla u^{(k+1)} - \mathbf{d}^{(k+1)} \\ b^{(k+1)} &= b^{(k)} + \int_{\Omega} u^{(k+1)} d\mathbf{x} - 1. \end{cases}$$

The shrink function is given by

$$\operatorname{shrink}(z, \eta) = \max\{|z| - \eta, 0\} \begin{pmatrix} z \\ |z| \end{pmatrix}.$$

Solving for $\mathbf{d}^{(k+1)}$ and $\mathbf{g}^{(k+1)}$ we use forward difference discretizations, namely

$$\nabla u^{(k+1)} = (u_{i+1,j} - u_{i,j}, u_{i,j+1} - u_{i,j})^T.$$

The Euler-Lagrange equations for the variable $u^{(k+1)}$ is

$$-\mu \frac{w(\mathbf{x})}{u(\mathbf{x})} + \lambda \operatorname{div} \theta - \alpha \left(\Delta u + \operatorname{div} \mathbf{g}^k - \operatorname{div} \mathbf{d}^k \right) + \gamma \left(\int_{\Omega} u \mathbf{x} + b_1^k - 1 \right) = 0.$$

Discretizing this simplifies to solving for the positive root of

$$(4 + \gamma) u_{i,j}^2 - \beta_{i,j} u_{i,j} - \mu w_{i,j} = 0$$

where

$$\begin{aligned} \beta_{i,j} &= u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1} - \lambda_2 \operatorname{div} \theta \\ &\quad - \alpha \left(d_{x,i,j}^k - d_{x,i-1,j}^k + d_{y,i,j}^k - d_{y,i,j-1}^k \right) \\ &\quad + \alpha \left(g_{x,i,j}^k - g_{x,i-1,j}^k + g_{y,i,j}^k - g_{y,i,j-1}^k \right) \\ &\quad + \gamma \left(1 - b^{(k)} - \sum_{(i',j') \neq (i,j)} u_{i',j'} \right). \end{aligned}$$

We solved for $u^{(k+1)}$ with a Gauss-Seidel solver. Heuristically, we found that using the relationships $\alpha = 2\mu N^2 M^2$ and $\gamma = 2\mu NM$ were sufficient for the solver to converge and provide good results. We also set λ to have values between 1.0 and 1.2. The parameter μ is the last remaining free parameter. This parameter can be chosen using V-cross validation or other techniques, such as the sparsity l_1 information criterion [17].

6 Results

In this section, we demonstrate the strengths of our models by providing several examples. We first show how our methods compare to existing methods for a dense data set. We then show our methods perform well for sparse data sets. Next, we explore an example with an aerial image and randomly selected events to show how these methods could be applied to geographic event data. Finally, we calculate probability density estimates for residential burglaries using our models.

6.1 Model Validation Example

To validate the use of our methods, we took a predefined probability map with sharp gradients that is shown in figure 2(a). The chosen valid region and the 8,000 selected events are displayed in figures 2(b) and 2(c), respectively. We performed density estimates with the Gaussian Kernel Density Estimate and the Total Variation MPLE method. The variance used for the Kernel Density Estimation is $\sigma = 2$. The results are provided in figures 2(d) and 2(e). The density estimates obtained from our Modified TV MPLE method and Weighted H_1 MPLE method are shown in figures 2(f) and 2(g), respectively. We also included our Weighted TV MPLE in figure 2(h).

Our methods maintain the boundary of the invalid region and appear close to the true solution. In addition, they keep the sharp gradient in the density estimate. The L_2 errors for these methods are located in Table 1.

	8,000 Events
Kernel Density Estimate	8.1079e-6
TV MPLE	6.6155e-6
Modified TV MPLE	4.1213e-6
Weighted H_1 MPLE	3.8775e-6
Weighted TV MPLE	4.3195e-6

Table 1: L_2 Error Comparison of Models for Dense Set of Events

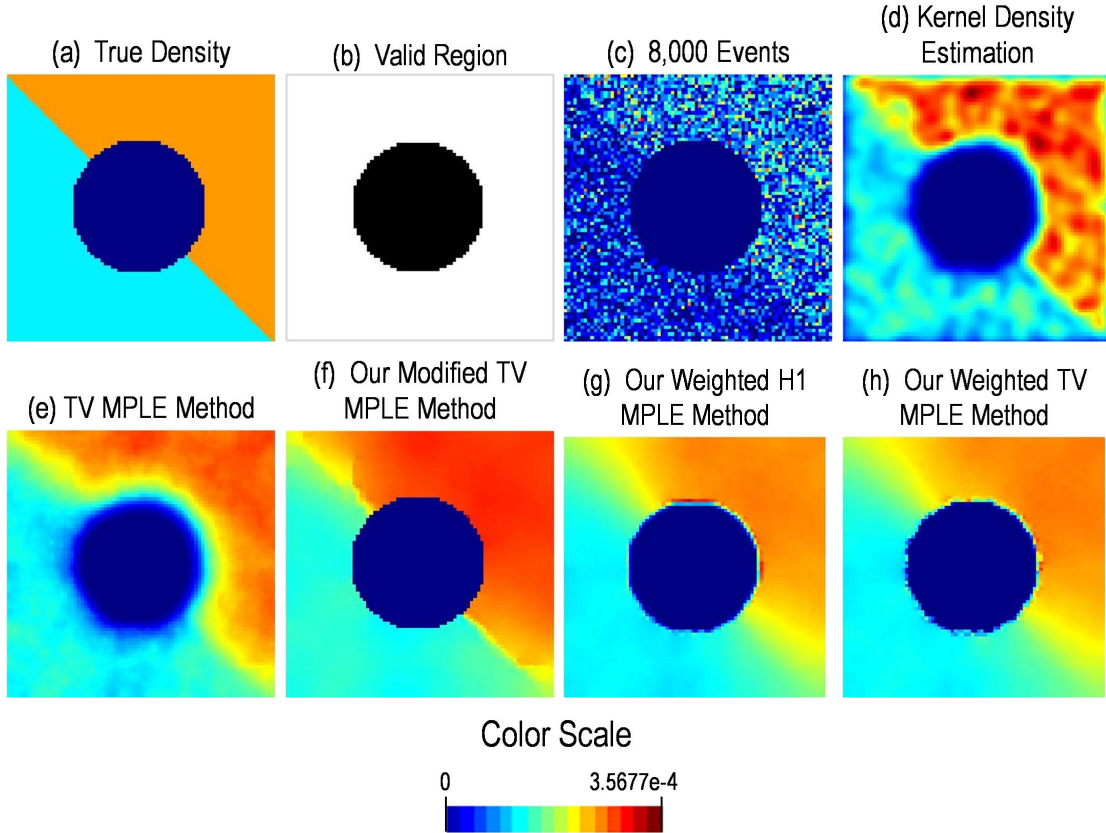


Figure 2: Model Validating Example with Dense Data Set

6.2 Sparse Data Example

Crimes and other types of events may be quite sparse in a given geographical region. Consequently, it becomes difficult to determine the probability that an event will occur in the area. It is challenging for density estimation methods that do not incorporate the spatial information to distinguish between invalid regions and areas that have not had any crimes but are still likely to have events. Using the same predefined probability density from the introduction section in figure 1(b), we demonstrate how our methods maintain these invalid regions for sparse data. The 40 events selected are shown in figure 3(b). The density estimates for current methods and our methods are given in figure 3. We used a variance $\sigma = 15$ for the Gaussian Kernel Density Estimate.

For this sparse problem, our Weighted H_1 MPLE and Modified TV MPLE methods maintain the boundary of the invalid region and appear close to the true solution. Table 2 contains the L_2 errors for both this example of 40 events and the example of 4,000 events from the introduction. Notice that our Modified TV and Weighted H_1 MPLE methods performed the best for both examples. The Weighted H_1 MPLE was exceptionally better for the sparse data set. The Weighted TV MPLE method does not perform as well for sparse data sets and fails to keep the boundary of the valid region. Since the rest of the examples contain sparse data sets, we will omit the Weighted TV MPLE method from the remaining sections.

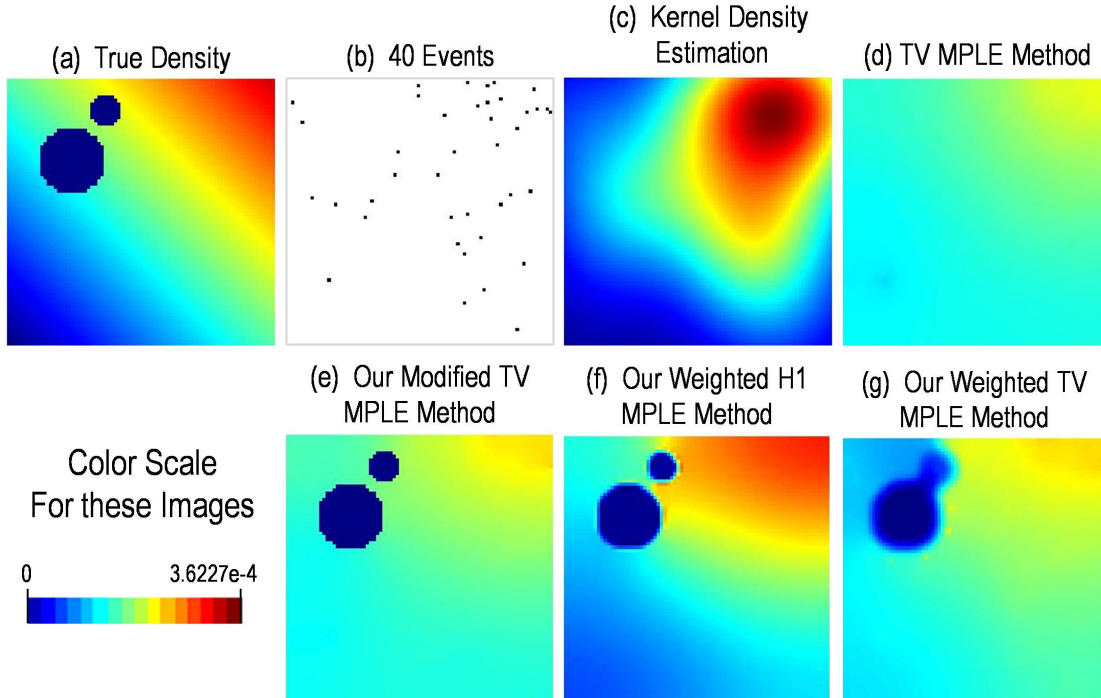


Figure 3: Sparse Example

	40 Events	4,000 Events
Kernel Density Estimate	2.3060e-5	7.3937e-6
TV MPLE	2.5347e-5	7.7628e-6
Modified TV MPLE	1.4345e-5	5.7996e-7
Weighted H_1 MPLE	3.8449e-6	2.1823e-6
Weighted TV MPLE	1.5982e-5	3.6179e-6

Table 2: L_2 Error Comparison of Models for Sparse Data Example and Introduction Example

6.3 Orange County Coastline Example

To test the models with external spatial data, we obtained from Google Earth TM a region of the Orange County coastline with clear invalid regions (see figure 4(a)). For the purposes of this example, it was determined to be impossible for events to occur in the ocean, rivers, or large parks located in the middle of the region. One may use various segmentation methods for selecting the valid region. For this example, we only have data from the true color aerial image, not multispectral data. To obtain the valid and invalid regions, we removed the “texture” (i.e. fine detailed features) using a Total Variation based denoising algorithm [16]. The resulting image, shown in figure 4(b), still contains detailed regions obtained from large features, such as large buildings. We wish to remove these and maintain prominent regional boundaries. Therefore, we smooth away from regions of large discontinuities. This is shown in figure 4(c). Since oceans, rivers, parks, and other such areas have generally lower intensity values than other regions, we threshold to find the boundary between the valid and invalid regions. The final valid region is displayed in figure 5(a).




Figure 4: Smoothing Regions of Orange County Coastline

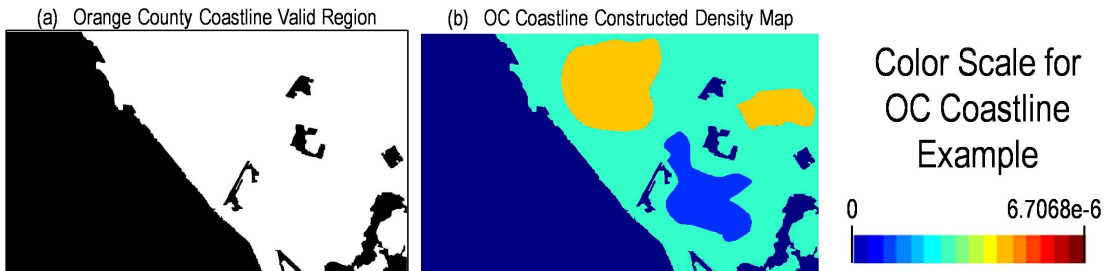


Figure 5: Orange County Coastline Valid Region and Constructed Density Map

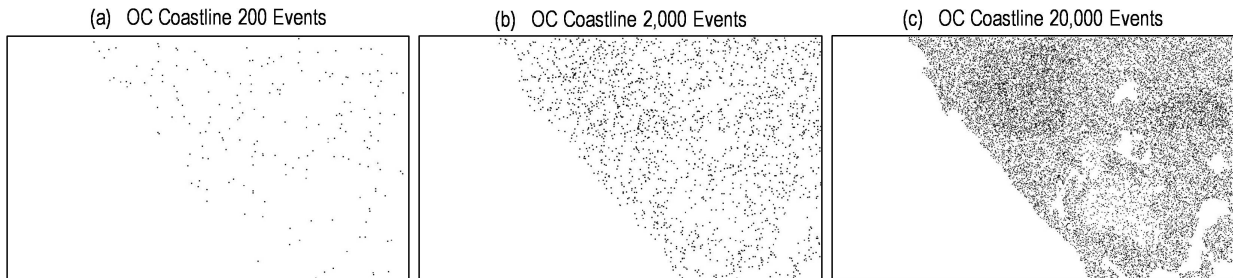


Figure 6: 200, 2,000, and 20,000 Sampled Events

From the valid region, we constructed a toy density map to represent the probability density for the example and to generate data. It is depicted in figure 5(b). Regions with colors farther to the right on the color scale are more likely to have events.

Sampling from this constructed density, we took distinct data sets of 200, 2,000 and 20,000 selected events given in figure 6. For each set of events, we included three probability density estimations for comparison. We first give the Gaussian Kernel Density Estimate followed by our Modified Total Variation MPLE model and our Weighted H_1 MPLE model. We provide all images together to allow for visual comparisons of the methods.

Summing up Gaussian distributions gives a smooth density estimate. Figure 7 contains the density estimates obtained using the Kernel Density Estimation model. The standard deviation σ of the Gaussians are given with each image. In all of these images, a nonzero density is estimated in the invalid region.

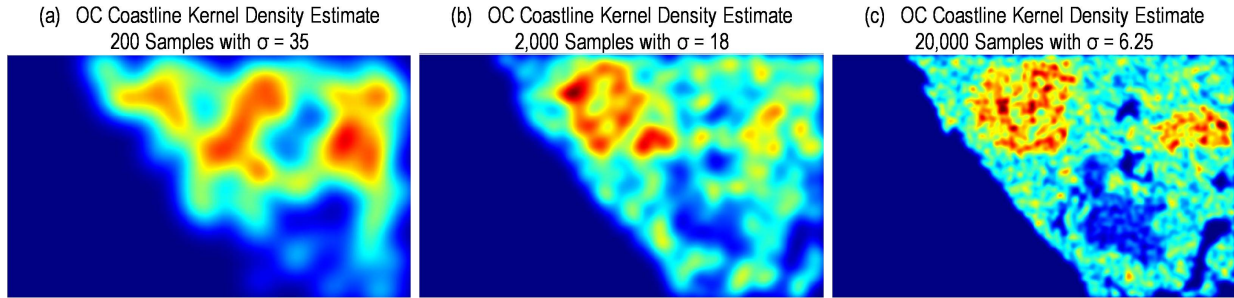


Figure 7: Kernel Density Estimates for 200, 2,000, and 20,000 Sampled Events

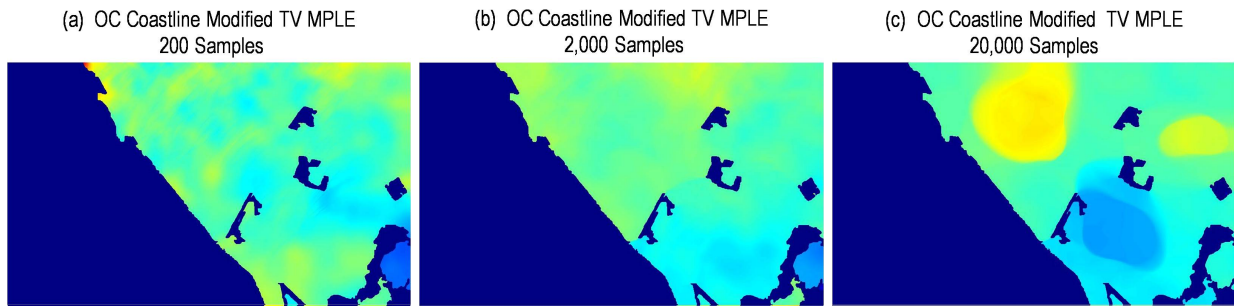


Figure 8: Modified TV MPLE for 200, 2,000, and 20,000 Sampled Events

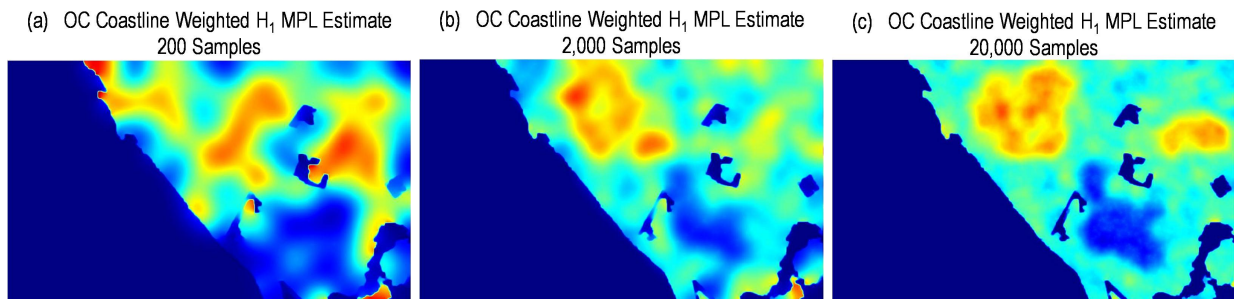


Figure 9: Weighted H_1 MPLE for 200, 2,000, and 20,000 Sampled Events

Taking the same set of events as the Kernel Density Estimation, the images in figure 8 were obtained using our first model, the Modified Total Variation MPLE method with the boundary edge aligning term. The parameter for λ must be sufficiently large in the TV method in order to prevent the diffusion of the density into the invalid region. In doing so, the boundary of the valid region may attain density values too large in comparison to the rest of the image when the size of the image is very large. To remedy this, we may take the resulting image from the algorithm and set the boundary of the valid region to zero and rescale the image to have a sum of one. The invalid region in this case sometimes has a very small non-zero estimate. For visualization purposes we have set this to zero. However, we note that the method has the strength that density does not diffuse through small sections of the invalid region back into the valid region on the opposite side. Events on one side of an object, such as a lake or river, should not necessarily predict events on the other side.

The next set of images in figure 9 estimate the density using the same sets of event data but with our Weighted H_1 MPLE model. Notice the difference for the invalid regions with our models and the Kernel Density Estimation model. This method does very well for the sparse data sets of 200 and 2,000 events.

6.3.1 Model Comparisons

The density estimates obtained from using our methods have a clear improvement in maintaining the boundary of the valid region. To determine how our models did in comparison to one another and to the Kernel Density Estimate, we calculated the L_2 errors located in Table 3. Our models consistently outperform the Kernel Density Estimation model. The Weighted H_1 MPLE method performs the best for the 2,000 and 20,000 events and visually appears closer to the true solution for the 200 events than the other methods. When the data is very sparse, the Modified TV MPLE method smooths everything to be approximately the same density, whereas the Weighted H_1 MPLE method keeps the density more localized.

	200 Events	2,000 Events	20,000 Events
Kernel Density Estimate	7.0338e-7	2.8847e-7	1.5825e-7
Modified TV MPLE	3.0796e-7	2.6594e-7	8.9353e-8
Weighted H_1 MPLE	5.4658e-7	1.5988e-7	5.8038e-8

Table 3: L_2 Error Comparison of Models for OC Coastline Example

6.4 Residential Burglary Example

The following example uses actual residential burglary information from the San Fernando Valley in Los Angeles. In figure 10 is the area of interest and the locations of 4,487 burglaries that occurred in the region during 2004 and 2005. The aerial image was obtained using Google EarthTM. We assume that residential burglaries cannot occur in large parks, lakes, mountainous areas without houses, at airports, and industrial areas. Using census or other types of data, housing density information for a given region can be calculated. Figure 10(c) is the housing density for our region of interest. The housing density provides us with the exact locations of where residential burglaries may occur. However, our methods prohibit the density estimates from spreading through the boundaries of the valid region. If we were to use this image directly as the valid region, then crimes on one side of a street will not have an effect on the opposite side of the road. Therefore, we fill in small holes and streets in the housing density image and use the image located in figure 10(d) as our valid region.

Using our Weighted H_1 MPLE and Modified TV MPLE models, the Gaussian Kernel Density Estimate with variance $\sigma = 21$, and the TV MPLE method, we obtained the density estimations shown in figure 11.

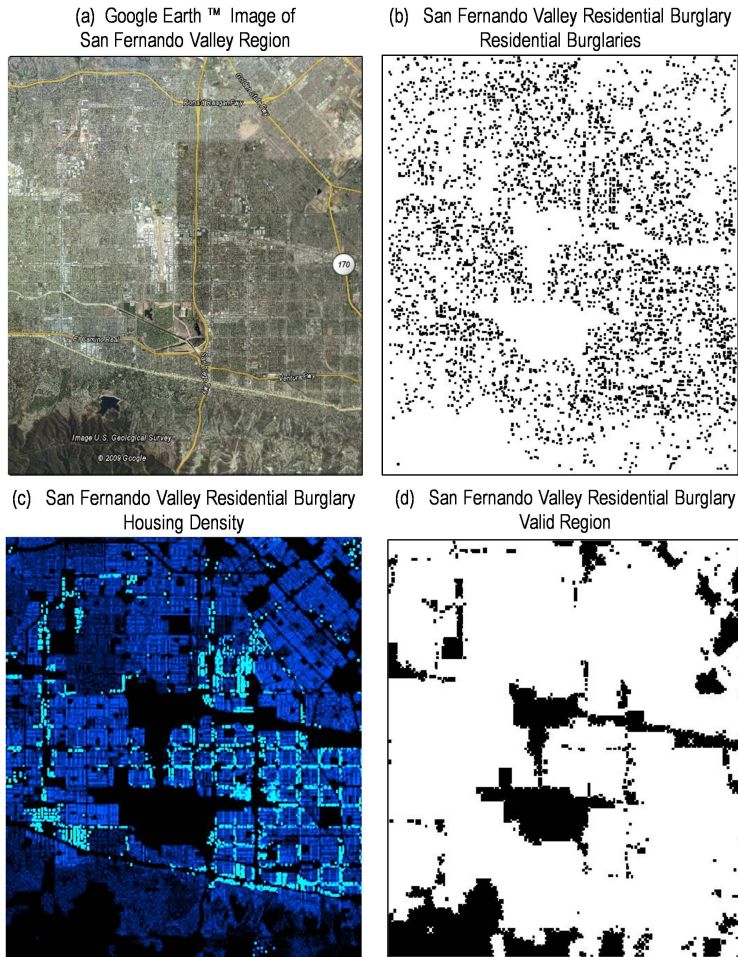


Figure 10: San Fernando Valley Aerial Image, its Residential Burglaries and Housing Density, and the Valid Region

7 Conclusions and Future Work

In this paper we have studied the problem of determining a more geographically accurate probability density estimate. We demonstrate the importance of this problem by showing how common density estimation techniques, such as Kernel Density Estimation, fail to restrict the support of the density in a set of realistic examples.

To handle this problem, we proposed a set of methods, based on Total Variation and H_1 -regularized MPLE models, that demonstrates great improvements in accurately enforcing the support of the density estimate when the valid region has been provided *a priori*. Unlike the TV-regularized methods, our H_1 model has the advantage that it performs well for very sparse data sets.

The effectiveness of the methods is shown in a set of examples in which burglary probability

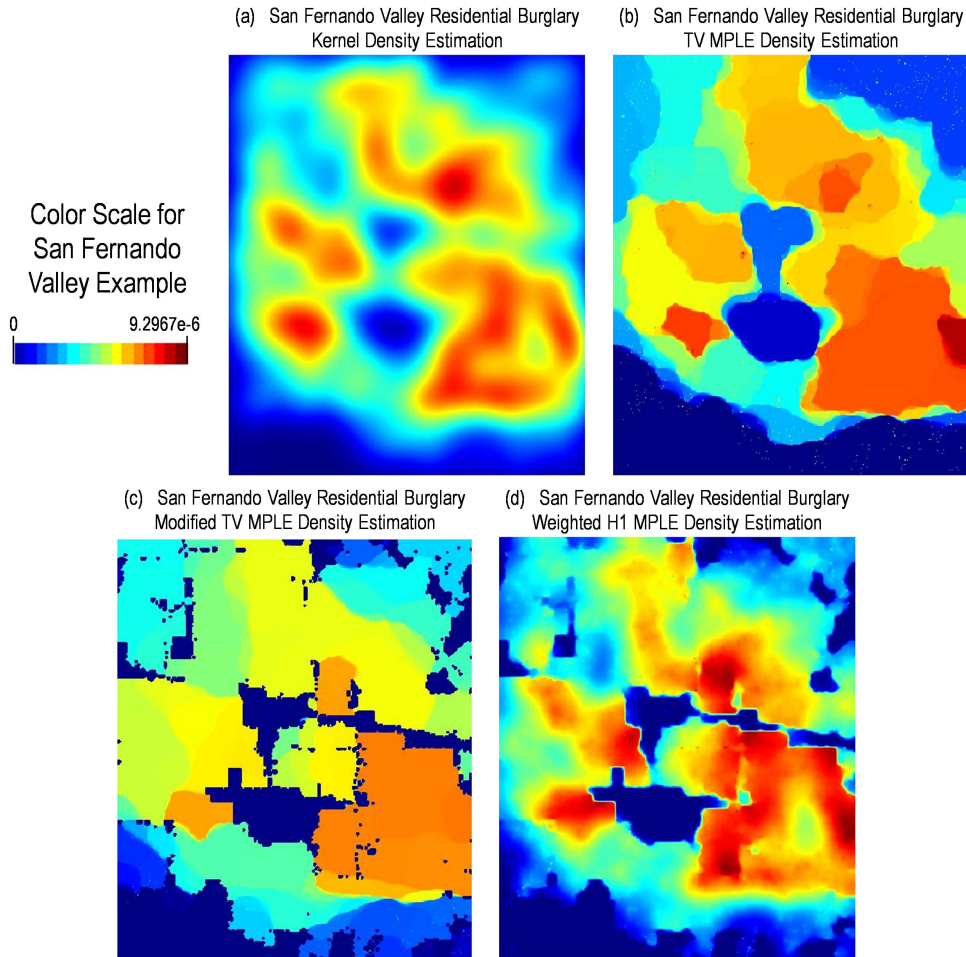


Figure 11: San Fernando Valley Density Estimates

densities are approximated from a set of crime events. Regions in which burglaries are impossible, such as oceans, mountains, and parks, are determined using aerial images or other external spatial data. These regions are then used to define an invalid region in which the density should be zero. Therefore, our methods are used to build geographically accurate probability maps.

It is interesting to note that there appears to be a relationship in the ratio between the number of samples and the size of the grid. In fact, each model has shown very different behavior in this respect. The TV based methods appear to be very sensitive to large changes in this ratio, whereas the H_1 method seems to be robust to these same changes. We are uncertain about why this phenomenon exists, and this would make an interesting future research topic.

There are many directions in which we can build on the results of this paper. We would like to devise better methods for determining the valid region, possibly evolving the edge set of the valid region using Γ -convergence [1]. Since this technique can be used for many types of event data, including residential burglaries, we would also like to apply this method to Iraq Body Count Data. Finally, we would like to handle possible errors in the data, such as incorrect positioning of events

that place them in the invalid region, by considering a probabilistic model of their position.

8 Acknowledgments

This work was supported by NSF grant BCS-0527388, NSF grant DMS-0914856, ARO MURI grant 50363-MA-MUR, ARO MURI grant W911NS-09-1-0559, ONR grant N000140810363 and the Department of Defense. We would like to thank George Tita and the LAPD for the burglary data set. We would also like to thank Jeff Brantingham, Martin Short, and the IPAM RIPS program at UCLA for the housing density data, which was obtained using ArcGIS and the LA County tax assessor data. We obtained our aerial images from Google Earth TM.

References

- [1] Luigi Ambrosio and Vincenzo Maria Tortorelli. Approximation of Functional Depending on Jumps by Elliptic Functional Via Γ -Convergence. *Communications on Pure and Applied Mathematics*, 43(8):999–1036, 1990.
- [2] Coloma Ballester, Vicent Caselles, Laura Igual, Joan Verdera, and Bernard Rougé. A Variational Model for P+XS Image Fusion. *Int. J. Comput. Vision*, 69(1):43–58, 2006.
- [3] P. L. Davies and A. Kovac. Densities, Spectral Densities and Modality. *The Annals of Statistics*, 32:1093–1136, 2004.
- [4] P. P. B. Eggermont and V. N. LaRiccia. *Maximum Penalized Likelihood Estimation*. Springer, 2001.
- [5] T. Goldstein and S. Osher. Split Bregman Method for L1 Regularized Problems. *SIAM Journal on Imaging Sciences*, 2:323–343, 2009.
- [6] I. J. Goodd and R. A. Gaskins. Nonparametric Roughness Penalties for Probability Densities. *Biometrika*, 58(2):255–277, 1971.
- [7] R. Koenker and I. Mizera. Density Estimation by Total Variation Regularization. *Advances in Statistical Modeling and Inference, Essays in Honor of Kjell A. Doksum*, World Scientific:613–634, 2007.
- [8] C. Kooperberg and C. J. Stone. Logspline Density Estimation with Free Knots. *Computational Statistics and Data Analysis*, 12:327–347, 2002.
- [9] Michael Moeller, Todd Wittman, and Andrea L. Bertozzi. Variational Wavelet Pan-Sharpener. Submitted to IEEE Transactions on Geoscience and Remote Sensing, December 2008.
- [10] George O. Mohler, Andrea L. Bertozzi, Thomas A. Goldstein, and Stanley J. Osher. Fast TV Regularization for 2D Maximum Penalized Likelihood Estimation. Technical Report 09-20, UCLA Department of Mathematics, 2009.

- [11] George O. Mohler and Martin B. Short. Geographic Profiling from Kinetic Models of Criminal Behavior. Preprint, 2009.
- [12] D. Mumford and J. Shah. Optimal Approximations by Piecewise Smooth Functions and Associated Variational Problems. *Communications on Pure and Applied Mathematics*, 42(5):577–685, 1989.
- [13] Mike O’Leary. The Mathematics of Geographic Profiling. *Journal of Investigative Psychology and Offender Profiling*, 6:253–265, 2009.
- [14] Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. An Iterative Regularization Method for Total Variation-Based Image Restoration. *Simul*, 4:460–489, 2005.
- [15] D. Kim Rossmo. *Geographic Profiling*. CRC Press, 2000.
- [16] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear Total Variation Based Noise Removal Algorithms. *Physica D*, 60:259–268, 1992.
- [17] S. Sardy and P. Tseng. Density Estimation by Total Variation Penalized Likelihood Driven by the Sparsity l_1 Information Criterion. To appear in the Scandinavian Journal of Statistics, 2009.
- [18] B. W. Silverman. Kernel Density Estimation Using the Fast Fourier Transform. *Statistical Algorithm, AS176, Appl. Statist.*, 31:93–97, 1982.
- [19] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, April 1986.
- [20] Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang. A New Alternating Minimization Algorithm for Total Variation Image Reconstruction. *SIAM Journal on Imaging Sciences*, 1(3):248–272, 2008.