Why Crime Doesn't Pay: Locating Criminals Through Geographic Profiling

Control Number: #7272

February 22, 2010

Abstract

Geographic profiling, the application of mathematics to criminology, has greatly improved police efforts to catch serial criminals by finding their residence. However, many geographic profiles either generate an extremely large area for police to cover or generates regions that are unstable with respect to internal parameters of the model. We propose, formulate, and test the Gaussian Rossmooth (GRS) Method, which takes the strongest elements from multiple existing methods and combines them into a more stable and robust model. We also propose and test a model to predict the location of the next crime. We tested our models on the Yorkshire Ripper case. Our results show that the GRS Method accurately predicts the location of the killer's residence. Additionally, the GRS Method is more stable with respect to internal parameters and more robust with respect to outliers than the existing methods. The model for predicting the location of the next crime generates a logical and reasonable region where the next crime may occur. We conclude that the GRS Method is a robust and stable model for creating a strong and effective model.

Contents

List of Figures

1 Introduction

Catching serial criminals is a daunting problem for law enforcement officers around the world. On the one hand, a limited amount of data is available to the police in terms of crimes scenes and witnesses. However, acquiring more data equates to waiting for another crime to be committed, which is an unacceptable trade-off. In this paper, we present a robust and stable **geographic profile** to predict the residence of the criminal and the possible locations of the next crime. Our model draws elements from multiple existing models and synthesizes them into a unified model that makes better use of certain empirical facts of criminology.

2 Plan of Attack

Our objective is to create a geographic profiling model that accurately describes the residence of the criminal and predicts possible locations for the next attack. In order to generate useful results, our model must incorporate two different schemes and must also describe possible locations of the next crime. Additionally, we must include assumptions and limitations of the model in order to ensure that it is used for maximum effectiveness.

To achieve this objective, we will proceed as follows:

1. **Define Terms -** This ensures that the reader understands what we are talking about and helps explain some of the assumptions and limitations of the model.
2. **Explain Existing Models -** This allows us to see how others have attacked the problem. Additionally, it provides a logical starting point for our model.
3. **Describe Properties of a Good Model-** This clarifies our objective and will generate a sketelon for our model.

With this underlying framework, we will present our model, test it with existing data, and compare it against other models.

3 Definitions

The following terms will be used throughout the paper:

1. **Spatial Mean -** Given a set of points, S, the spatial mean is the point that represents the middle of the data set.
2. **Standard Distance -** The standard distance is the analog of standard deviation for the spatial mean.

3. **Marauder -** A serial criminal whose crimes are situated around his or her place of residence.

4. **Distance Decay -** An empirical phenomenon where criminal don't travel too far to commit their crimes.

5. **Buffer Area -** A region around the criminal's residence or workplace where he or she does not commit crimes.[1] There is some dispute as to whether this region exists. [2] In our model, we assume that the buffer area exists and we measure it in the same spatial unit used to describe the relative locations of other crime scenes.

6. **Manhattan Distance -** Given points $\mathbf{a} = (x_1, y_1)$ and $\mathbf{b} = (x_2, y_2)$, the Manhattan distance from **a** to **b** is $|x_1 - x_2| + |y_1 - y_2|$. This is also known as the $\mathbf{1 - norm}$.

7. **Nearest Neighbor Distance -** Given a set of points S, the nearest neighbor distance for a point $x \in S$ is

$$\min_{s \in S - \{x\}} |x - s|$$

   Any norm can be chosen.

8. **Hot Zone -** A region where a predictive model states that a criminal might be. Hot zones have much higher predictive scores than other regions of the map.

9. **Cold Zone -** A region where a predictive model scores exceptionally low.

4 Existing Methods

Currently there are several existing methods for interpolating the position of a criminal given the location of the crimes.

4.1 Great Circle Method

In the **great circle method**, the distances between crimes are computed and the two most distant crimes are chosen. Then, a great circle is drawn so that both of the points are on the great circle. The midpoint of this great circle is then the assumed location of the criminal's residence and the area bounded by the great circle is where the criminal operates. This model is computationally inexpensive and easy to understand. [3] Moreover, it is easy to use and requires very little training in order to master the technique.[2] However, it has certain drawbacks. For example, the area given by this method is often very large and other studies have shown that a smaller area suffices. [4] Additionally, a few outliers can generate an even larger search area, thereby further slowing the police effort.

4.2 Centrography

In **centrography**, crimes are assigned x and y coordinates and the "center of mass" is computed as follows:

$$x_{center} = \sum_{i=1}^{n} \frac{x_i}{n}$$

$$y_{center} = \sum_{i=1}^{n} \frac{y_i}{n}$$

Intuitively, centrography finds the mean $x-$coordinate and the mean y-coordinate and associates this pair with the criminal's residence (this is called the **spatial mean**). However, this method has several flaws. First, it can be unstable with respect to outliers. Consider the following set of points (shown in Figure 1:

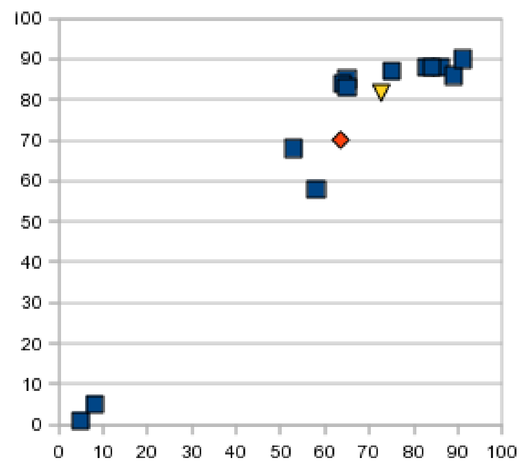

Figure 1: The effect of outliers upon centrography. The current spatial mean is at the red diamond. If the two outliers in the lower left corner were removed, then the center of mass would be located at the yellow triangle.

Though several of the crime scenes (blue points) in this example are located in a pair of upper clusters, the spatial mean (red point) is reasonably far away from the clusters. If the two outliers are removed, then the spatial mean (yellow point) is located closer to the two clusters.

A similar method uses the median of the points. The median is not so strongly affected by outliers and hence is a more stable measure of the middle.[3]

Alternatively, we can circumvent the stability problem by incorporating the 2-D analog of standard deviation called the **standard distance**:

$$\sigma_{SD} = \sqrt{\frac{\sum d_{center,i}}{N}}$$

where N is the number of crimes committed and $d_{center,i}$ is the distance from the spatial center to the i^{th} crime.

By incorporating the standard distance, we get an idea of how "close together" the data is. If the standard distance is small, then the kills are close together. However, if the standard distance is large, then the kills are far apart.

Unfortunately, this leads to another problem. Consider the following data set (shown in Figure 2):

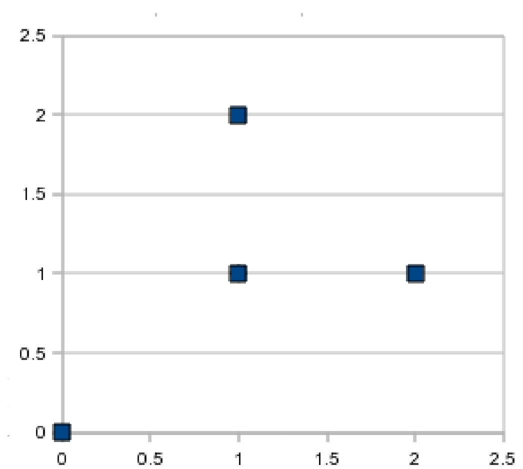

Figure 2: Crimes scenes that are located very close together can yield illogical results for the spatial mean. In this image, the spatial mean is located at the same point as one of the crime scenes at (1,1).

In this example, the kills (blue) are closely clustered together, which means that the centrography model will yield a center of mass that is in the middle of these crimes (in this case, the spatial mean is located at the same point as one of the crimes). This is a somewhat paradoxical result as research in criminology suggests that there is a **buffer area** around a serial criminal's place of residence where he or she avoids the commission of crimes.[3, 1] That is, the potential kill area is an annulus. This leads to **Rossmo's formula**[1], another mathematical model that predicts the location of a criminal.

4.3 Rossmo's Formula

Rossmo's formula divides the map of a crime scene into grid with i rows and j columns. Then, the probability that the criminal is located in the box at row i and column j is

$$P_{i,j} = k \sum_{c=1}^{T} \left[\frac{\phi}{(|x_i - x_c| + |y_j - y_c|)^f} + \frac{(1-\phi)(B^{g-f})}{(2B - |x_i - x_c| - |y_j - y_c|)^g} \right]$$

where $f = g = 1.2$, k is a scaling constant (so that P is a probability function), T is the total number of crimes, ϕ puts more weight on one metric than the other, and B is the radius of the buffer zone (and is suggested to be one-half the mean of the **nearest neighbor distance** between crimes).[1] Rossmo's formula incorporates two important ideas:

1. Criminals won't travel too far to commit their crimes. This is known as **distance decay**.
2. There is a buffer area around the criminal's residence where the crimes are less likely to be committed.

However, Rossmo's formula has two drawbacks. If for any crime scene x_c, y_c, the equality $2B = |x_i - x_c| + |y_j - y_c|$, is satisfied, then the term $\frac{(1-\phi)(B^{g-f})}{(2B - |x_i - x_c| - |y_j - y_c|)^g}$ is undefined, as the denominator is 0. Additionally, if the region associated with ij is the same region as the crime scene, then $\frac{\phi}{(|x_i - x_c| + |y_j - y_c|)^f}$ is undefined by the same reasoning. Figure 3 illustrates this:

This "delta function-like" behavior is disconcerting as it essentially states that the criminal either lives right next to the crime scene or on the boundary defined by Rossmo. Hence, the B-value becomes exceptionally important and needs its own heuristic to ensure its accuracy. A non-optimal choice of B can result in highly unstable search zones that vary when B is altered slightly.

5 Assumptions

Our model is an expansion and adjustment of two existing models, centrography and Rossmo's formula, which have their own underlying assumptions. In order to create an effective model, we will make the following assumptions:

1. ***The buffer area exists*** - This is a necessary assumption and is the basis for one of the mathematical components of our model.
2. ***More than 5 crimes have occurred*** - This assumption is important as it ensures that we have enough data to make an accurate model. Additionally, Rossmo's model stipulates that 5 crimes have occurred[1].

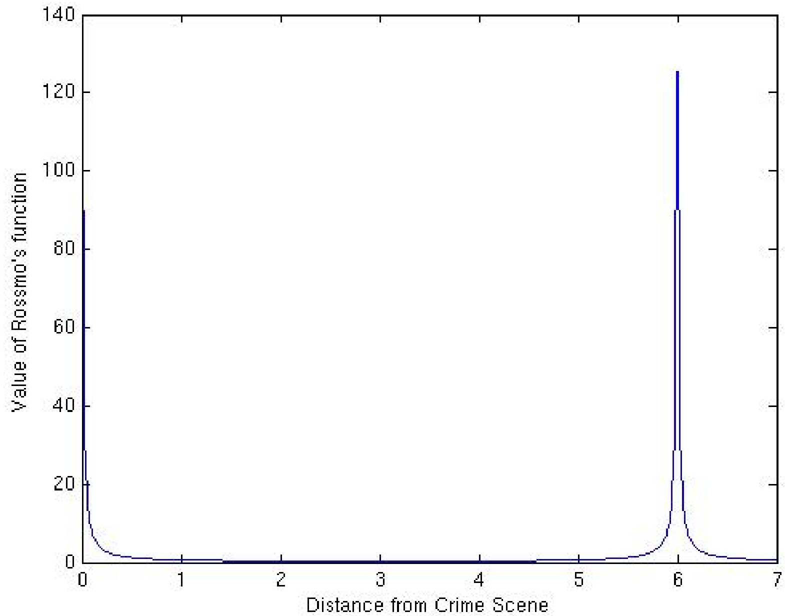

Figure 3: The summand in Rossmo's formula (2B = 6). Note that the function is essentially 0 at all points except for the scene of the crime and at the buffer zone and is undefined at those points

3. ***The criminal only resides in one location*** - By this, we mean that though the criminal may change residence, he or she will not move to a completely different area and commit crimes there. Empirically, this assumption holds, with a few exceptions such as David Berkowitz[1]. The importance of this assumption is it allows us to adapt Rossmo's formula and the centrography model. Both of these models implicitly assume that the criminal resides in only one general location and is not nomadic.

4. ***The criminal is a marauder*** - This assumption is implicitly made by Rossmo's model as his spatial partition method only considers a small rectangular region that contains all of the crimes.

With these assumptions, we present our model, the Gaussian Rossmooth method.

6 Gaussian Rossmooth

6.1 Properties of a Good Model

Much of the literature regarding criminology and geographic profiling contains criticism of existing models for catching criminals. [1, 2] From these criticisms, we develop the following criteria for creating a good model:

1. **Gives an accurate prediction for the location of the criminal** - This is vital as the objective of this model is to locate the serial criminal. Obviously, the model cannot give a definite location of the criminal, but it should at least give law enforcement officials a good idea where to look.

2. **Provides a good estimate of the location of the next crime** - This objective is slightly harder than the first one, as the criminal can choose the location of the next crime. Nonetheless, our model should generate a region where law enforcement can work to prevent the next crime.

3. **Robust with respect to outliers** - Outliers can severely skew predictions such as the one from the centrography model. A good model will be able to identify outliers and prevent them from adversely affecting the computation.

4. **Consitent within a given data set** - That is, if we eliminate data points from the set, they do not cause the estimation of the criminal's location to change excessively. Additionally, we note that if there are, for example, eight murders by one serial killer, then our model should give a similar prediction of the killer's residence when it considers the first five, first six, first seven, and all eight murders.

5. **Easy to compute** - We want a model that does not entail excessive computation time. Hence, law enforcement will be able to get their information more quickly and proceed with the case.

6. **Takes into account empirical trends** - There is a vast amount of empirical data regarding serial criminals and how they operate. A good model will incorporate this data in order to minimize the necessary search area.

7. **Tolerates changes in internal parameters** - When we tested Rossmo's formula, we found that it was not very tolerant to changes of the internal parameters. For example, varying B resulted in substantial changes in the search area. Our model should be stable with respect to its parameters, meaning that a small change in any parameter should result in a small change in the search area.

6.2 Outline of Our Model

We know that centrography and Rossmo's method can both yield valuable results. When we used the mean and the median to calculate the centroid of a string of murders in Yorkshire, England, we found that both the median-based and mean-based centroid were located very close to the home of the criminal. Additionally, Rossmo's method is famous for having predicted the home of a criminal in Louisiana. In our approach to this problem, we adapt these methods to preserve their strengths while mitigating their weaknesses.

1. **Smoothen Rossmo's formula** - While the theory behind Rossmo's formula is well documented, its implementation is flawed in that his formula reaches asymptotes when the distance away from a crime scene is 0 (i.e. point (x_i, y_j) is a crime scene), or when a point is exactly $2B$ away from a crime scene. We must smoothen Rossmo's formula so that idea of a buffer area is mantained, but the asymptotic behavior is removed and the tolerance for error is increased.

2. **Incorporate the spatial mean** - Using the existing crime scenes, we will compute the spatial mean. Then, we will insert a Gaussian distribution centered at that point on the map. Hence, areas near the spatial mean are more likely to come up as hot zones while areas further away from the spatial mean are less likely to be viewed as hot zones. This ensures that the intuitive idea of centrography is incorporated in the model and also provides a general area to search. Moreover, it mitigates the effect of outliers by giving a probability boost to regions close to the center of mass, meaning that outliers are unlikely to show up as hot zones.

3. **Place more weight on the first crime** - Research indicates that criminals tend to commit their first crime closer to their home than their latter ones.[5] By placing more weight on the first crime, we can create a model that more effectively utilizes criminal psychology and statistics.

6.3 Our Method

6.3.1 Rossmooth Method

First, we eliminated the scaling constant k in Rossmo's equation. As such, the function is no longer a probability function but shows the relative likelihood of the criminal living in a certain sector. In order to eliminate the various spikes in Rossmo's method, we altered the distance decay function.

We wanted a distance decay function that:

1. Preserved the distance decay effect. Mathematically, this meant that the function decreased to 0 as the distance tended to infinity.

2. Had an interval around the buffer area where the function values were close to each other. Therefore, the criminal could ostensibly live in a small region around the buffer zone, which would increase the tolerance of the B-value.

We examined various distance decay functions [1, 3] and found that the functions resembled $f(x) = Ce^{-m(x-x_0)^2}$. Hence, we replaced the second term in Rossmo's function with term of the form $(1-\phi) \times Ce^{-k(x-x_0)^2}$. Our modified equation was:

$$E_{i,j} = \sum_{c=1}^{T} \left[\frac{\phi}{(|x_i - x_c| + |y_j - y_c|)^f} + (1-\phi) \times Ce^{-(2B-(|x_i-x_c|+|y_j-y_c|))^2} \right]$$

However, this maintained the problematic region around any crime scene. In order to eliminate this problem, we set an `EPSILON` so that any point within `EPSILON` (defined to be 0.5 spatial units) of a crime scene would have a weighting of a constant cap. This prevented the function from reaching an asymptote as it did in Rossmo's model. The cap was defined as

$$\text{CAP} = \frac{\phi}{\texttt{EPSILON}^f}$$

The C in our modified Rossmo's function was also set to this cap. This way, the two maximums of our modified Rossmo's function would be equal and would be located at the crime scene and the buffer zone.

This function yielded the following curve (shown in in Figure 4), which fit both of our criteria:

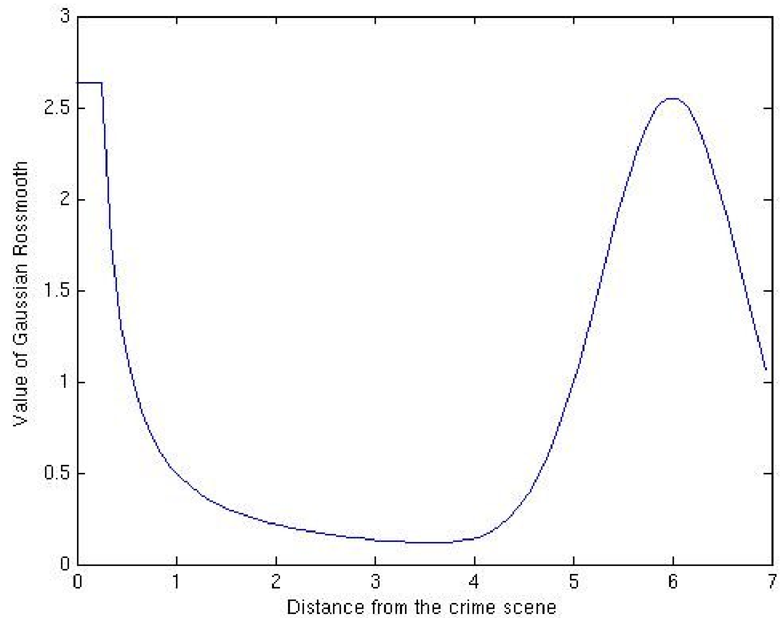

Figure 4: The summand in smoothed Rossmo's formula ($2B = 6$, $\phi = 0.5$, and $\texttt{EPSILON} = 0.5$). Note that there is now a region around the buffer zone where the value of the function no longer changes very rapidly.

At this point, we noted that E_{ij} had served its purpose and could be replaced in order to create a more intuitive idea of how the function works. Hence, we replaced $E_{i,j}$ with the following sum:

$$\sum_{c=1}^{T} [D_1(c) + D_2(c)]$$

where:

$$D_1(c) = \min\left(\frac{\phi}{(|x_i - x_c| + |y_j - y_c|)^f} \ , \ \frac{\phi}{\texttt{EPSILON}^f} \right)$$

$$D_2(c) = (1 - \phi) \times Ce^{-(2B-(|x_i-x_c|+|y_j-y_c|))^2}$$

For equal weighting on both $D_1(c)$ and $D_2(c)$, we set ϕ to 0.5.

6.3.2 Gaussian Rossmooth Method

Now, in order to incorporate the inuitive method, we used centrography to locate the center of mass. Then, we generated a Gaussian function centered at this point. The Gaussian was given by:

$$G = Ae^{-\left(\frac{(x-x_{center})^2}{2\sigma_x^2}+\frac{(y-y_{center})^2}{2\sigma_y^2}\right)}$$

where A is the amplitude of the peak of the Gaussian. We determined that the optimal A was equal to 2 times the cap defined in our modified Rossmo's equation. ($A = \frac{2\phi}{\texttt{EPSILON}^f}$)

To deal with empirical evidence that the first crime was usually the closest to the criminal's residence, we doubled the weighting on the first crime. However, the weighting can be represented by a constant, W. Hence, our final Gaussian Rosmooth function was:

$$GRS(x_i, y_j) = G + W(D_1(1) + D_2(1)) + \sum_{c=2}^{T} [D_1(c) + D_2(c)]$$

7 Gaussian Rossmooth in Action

7.1 Four Corners: A Simple Test Case

In order to test our Gaussain Rossmooth (GRS) method, we tried it against a very simple test case. We placed crimes on the four corners of a square. Then, we hypothesized that the model would predict the criminal to live in the center of the grid, with a slightly higher hot zone targeted toward the location of the first crime. Figure 5 shows our results, which fits our hypothesis.

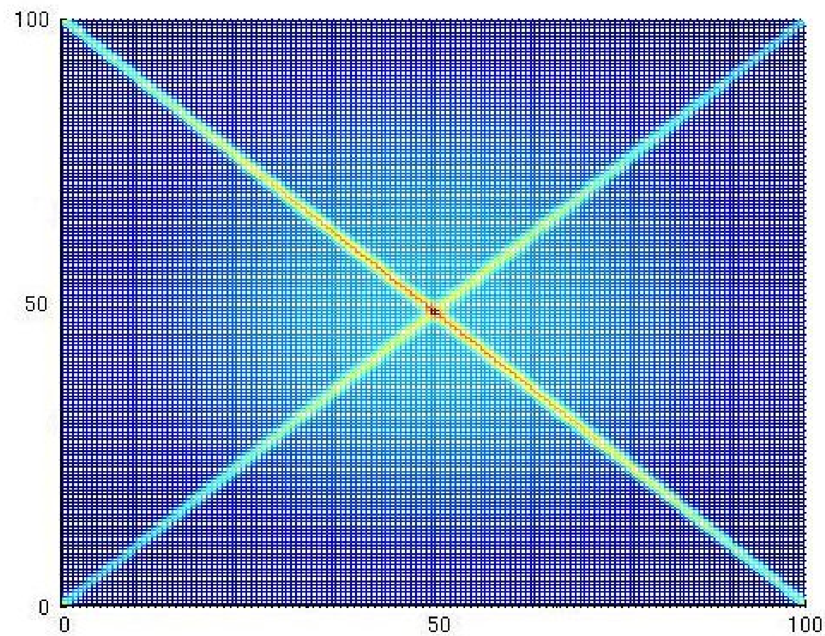

Figure 5: The Four Corners Test Case. Note that the highest hot spot is located at the center of the grid, just as the mathematics indicates.

7.2 Yorkshire Ripper: A Real-World Application of the GRS Method

After the model passed a simple test case, we entered the data from the Yorkshire Ripper case. The Yorkshire Ripper (a.k.a. Peter Sutcliffe) committed a string of 13 murders and several assaults around Northern England. Figure 6 shows the crimes of the Yorkshire Ripper and the locations of his residence[1] :

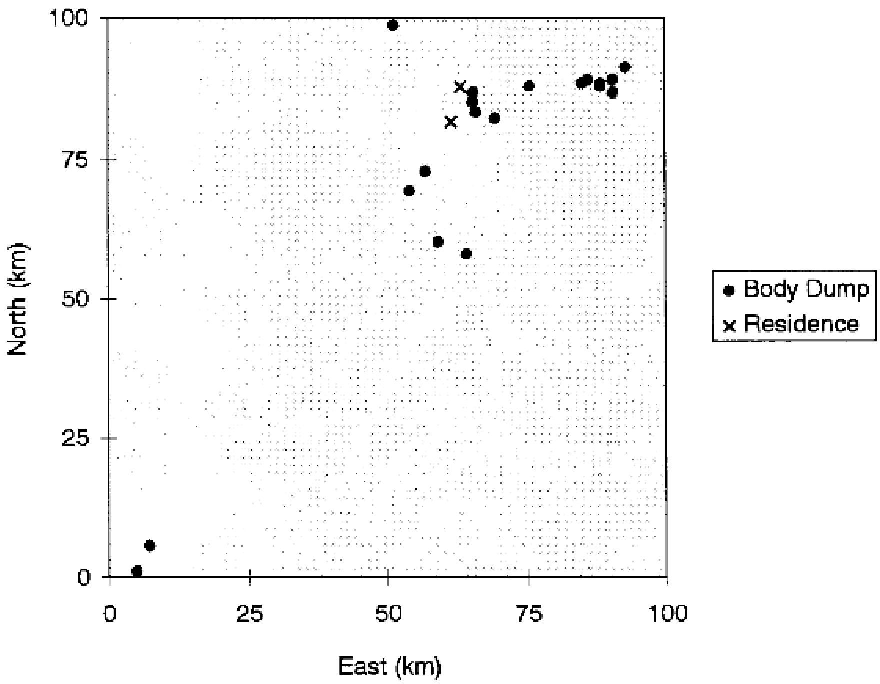

Figure 6: Crimes and residences of the Yorkshire Ripper. There are two residences as the Ripper moved in the middle of the case. Some of the crime locations are assaults and others are murders.

When our full model ran on the murder locations, our data yielded the image show in Figure 7:

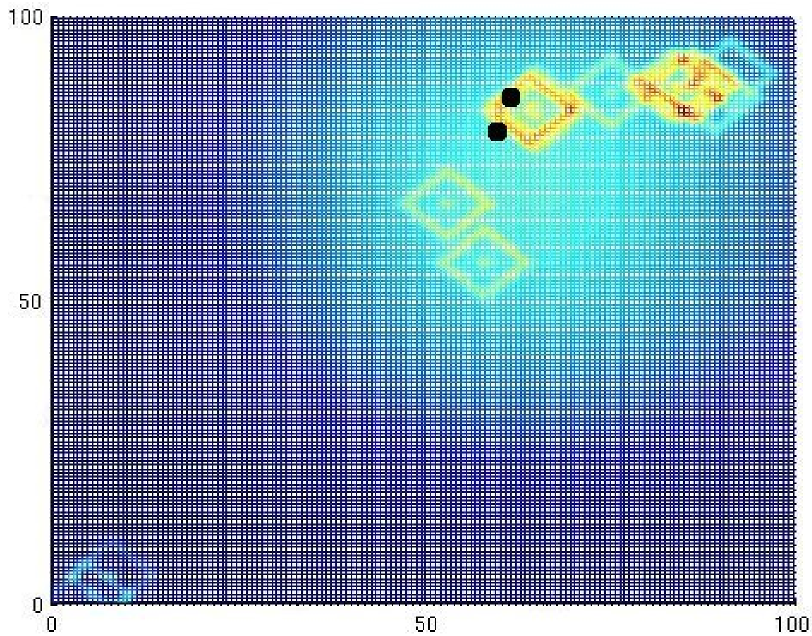

Figure 7: GRS output for the Yorkshire Ripper case (B = 2.846). Black dots indicate the two residences of the killer.

In this image, hot zones are in red, orange, or yellow while cold zones are in black and blue. Note that the Ripper's two residences are located in the vicinity of our hot zones, which shows that our model is at least somewhat accurate. Additionally, regions far away from the center of mass are also blue and black, regardless of whether a kill happened there or not.

7.3 Sensitivity Analysis of Gaussian Rossmooth

The GRS method was exceptionally stable with respect to the parameter B. When we ran Rossmo's model, we found that slight variations in B could create drastic variations in the given distribution. On many occassions, a change of 1 spatial unit in B caused Rossmo's method to destroy high value regions and replace them with mid-level value or low value regions (i.e., the region would completely dissapper). By contrast, our GRS method scaled the hot zones.

Figures 8 and 9 show runs of the Yorkshire Ripper case with B-values of 2 and 4 respectively. The black dots again correspond to the residence of the criminal. The original run (Figure 7) had a B-value of 2.846. The original B-value was obtained by using Rossmo's nearest neighbor distance metric. Note that when B is varied, the size of the hot zone varies, but the shape of the hot zone does not. Additionally, note that when a B-value gets further away from the value obtained by the nearest neighbor distance metric, the accuracy of the model decreases slightly, but the overall search areas are still quite accurate.

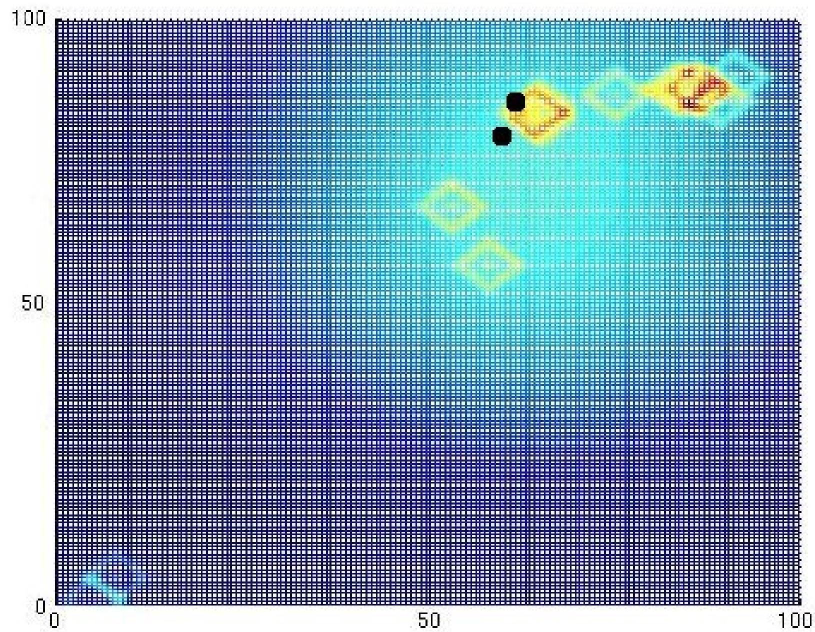

Figure 8: GRS method run on Yorkshire Ripper data ($B = 2$). Note that the major difference between this model and Figure 7 is that the hot zones in this figure are smaller than in the original run.

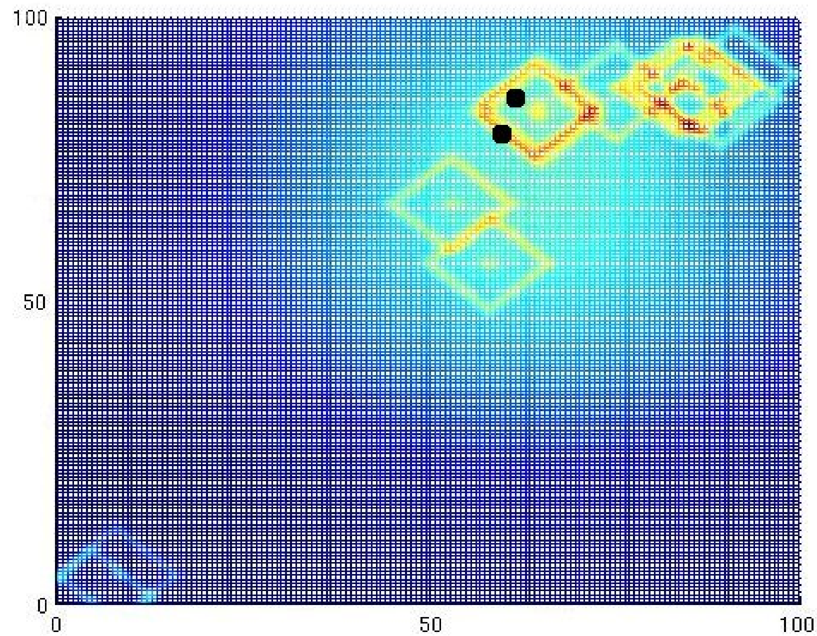

Figure 9: GRS method run on Yorkshire Ripper data ($B = 4$). Note that the major difference between this model and Figure 7 is that the hot zones in this figure are larger than in the original run.

7.4 Self-Consistency of Gaussian Rossmooth

In order to test the self-consistency of the GRS method, we ran the model on the first N kills from the Yorkshire Ripper data, where N ranged from 6 to 13, inclusive. The self-consistency of the GRS method was adversely affected by the center of mass correction, but as the case number approached 11, the model stabilized. This phenomenon can also be attributed to the fact that the Yorkshire Ripper's crimes were more separated than those of most marauders. A selection of these images can be viewed in the appendix.

8 Predicting the Next Crime

The GRS method generates a set of possible locations for the criminal's residence. We will now present two possible methods for predicting the location of the criminal's next attack. One method is computationally expensive, but more rigorous while the other method is computationally inexpensive, but more intuitive.

8.1 Matrix Method

Given the parameters of the GRS method, the region analyzed will be a square with side length n spatial units. Then, the output from the GRS method can be interpreted as an $n \times n$ matrix. Hence, for any two runs, we can take the norm of their matrix difference and compare how similar the runs were. With this in mind, we generate the following method.

For every point on the grid:

1. Add crime to this point on the grid.
2. Run the GRS method with the new set of crime points.
3. Compare the matrix generated with these points to the original matrix by subtracting the components of the original matrix from the components of the new matrix.
4. Take a **matrix norm** of this difference matrix.
5. Remove the crime from this point on the grid.

As a lower matrix norm indicates a matrix similar to our original run, we seek the points so that the matrix norm is minimized.

There are several matrix norms to choose from. We chose the **Frobenius norm** because it takes into account all points on the difference matrix.[6] The Frobenius norm is:

$$||A||_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n}|a_{ij}|^2}$$

However, the Matrix Method has one serious drawback: it is exceptionally expensive to compute. Given an $n \times n$ matrix of points and c crimes, the GRS method runs in $O(cn^2)$. As the Matrix method runs the GRS method at each of n^2 points, we see that the Matrix Method runs in $O(cn^4)$. With the Yorkshire Ripper case, $c = 13$ and $n = 151$. Accordingly, it requires a fairly long time to predict the location of the next crime. Hence, we present an alternative solution that is more intuitive and efficient.

8.2 Boundary Method

The Boundary Method searches the GRS output for the highest point. Then, it computes the average distance, r, from this point to the crime scenes. In order to generate a resonable search area, it discards all outliers (i.e., points that were several times further away from the high point than the rest of the crimes scenes.) Then, it draws annuli of outer radius r (in the 1-norm sense) around all points above a certain cutoff value, defined to be 60% of the maximum value. This value was chosen as it was a high enough percentage value to contain all of the hot zones.

The beauty of this method is that essentially it uses the same algorithm as the GRS. We take all points on the hot zone and set them to "crime scenes." Recall that our GRS formula was :

$$GRS(x_i, y_j) = G + W(D_1(1) + D_2(1)) + \sum_{c=2}^{T} [(D_1(c) + D_2(c))]$$

In our boundary model, we only take the terms that involve $D_2(c)$. However, let $D_2'(c)$ be a modified $D_2(c)$ defined as follows:

$$D_2'(c) = (1 - \phi) \times Ce^{-(r-(|x_i - x_c| + |y_j - y_c|))^2}$$

Then, the boundary model is:

$$BS(x_i, y_j) = \sum_{c=1}^{T} D_2'(c)$$

9 Boundary Method in Action

This model generates an outer boundary for the criminal's next crime. However, our model does not fill in the region within the inner boundary of the annulus. This region should still be searched as the criminal may commit crimes here. Figure 10 shows the boundary generated by analyzing the Yorkshire Ripper case.

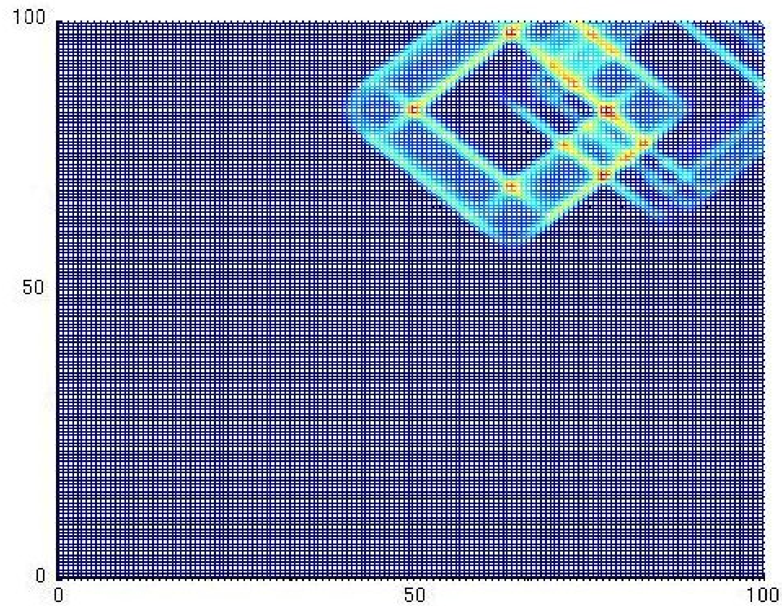

Figure 10: The boundary region generated by our Boundary Method. Note that boundary region covers many of the crimes committed by the Sutcliffe.

Even more astonishingly, the boundary method almost perfectly describes the location of the 12^{th} crime when it is run on the first 11 crimes. The table below shows the highest value predicted locations of crimes scenes and their actual locations:

Crime #	Actual Crime	Predicted Crime	Distance Difference
11	(65, 83)	(72, 88)	12
12	(75, 87)	(72, 88)	4
13	(84, 88)	(78, 84)	10

10 Limitations

Like any predictive model, our model has limitations.

1. **There are exceptions to our weighting heuristic** - By this, we mean that some criminals make their first strike far away from home. For example, Peter Sutcliffe's first victim was located over 13 miles away from his place of residence, much further away from his home than some of his

later crimes.[1] Hence, though our model may predict the right location, it can also place hot zones in regions where the criminal does not live.

2. **Our model only deals with criminal who are marauders** - Some criminals commit their crimes in a region far away from their place of residence. These criminals, known as **commuters**, will not be so effectively tracked by our model. However, when we generated the model, we made the explict assumption that the criminal was a marauder, so if law enforcement has evidence that the criminal is a commuter, then they should keep that in mind while using our model. (However, Peter Sutcliffe was on the border of commuter and marauder[1], meaning that our model can still give a surprisingly accurate result.)

11 Executive Summary

Our approach to generating a geographic profile was to combine two existing models. We ran our model on test data from the Yorkshire Ripper case and it yielded more precise and accurate estimates of the killer's two residences than other existing models. Therefore, our model has been shown to work on real world data.

11.1 Outline of Our Model

Our model finds points that allow the criminal to reach all crime scenes with the lowest possible total effort. It then factors in the buffer area effect, the tendency of criminals to avoid committing crimes too close to home. Then, our model generates "hot zones" where the criminal may live. These areas are where police efforts should start. It also ranks all regions thereby showing were police efforts are probably unnecessary and can be diverted toward other regions.

11.2 Running the Model

Our model requires the input of the locations of crimes. A grid is also required so that the crimes can be plotted on that grid. Our model contains several internal parameters that can be changed to sharpen the model. Most of these parameters are set to values that do not need to change from case to case. However, a very important parameter is the buffer area. In order to figure out the buffer area, find one-half of the average of the nearest neighbor distance of the crime scenes. In general, this turns out to be the radius of the buffer area (for murders). For other types of crimes, the B value is smaller.

11.3 Interpreting the Results

Our model generates "hot zones", regions where the criminal may live. The hot zone are ranked on the following scale (from highest to lowest):

1. **Red**
2. **Orange**
3. **Yellow**
4. **Green**
5. **Blue**

The higher the value of the zone, the more priority should be given to searching this region. Regions with rankings from Red to Yellow should be searched while regions at or below Green can be avoided until the other regions have been searched.

In the model for predicting the next crime, we generate a region where the next crime is likely to occur. This calculation is made by looking at hot zones and factoring in how far the criminal's residence is from the location. The images created by this model map the outer boundary where a crime is likely to occur. However, the region inside of this boundary should also be searched because the criminal may change his or her actions due to the increased police presence.

11.4 Limitations

Our model applies to cases where at least 5 crimes have occurred. The model can still be used with fewer than 5 crimes, but it will not be as accurate. Additionally, our model should be applied when law enforcement believes that the criminal lives inside of the area bounded by the crimes or near the area where the crimes happened. If the criminal lives outside of the region bounded by the kills, then our model may generate an inaccurate profile. Furthermore, our model doesn't take terrain into account. When using the model, a human touch is need in order to determine if the model is providing reasonable hot zones. If a hot zone appears in the middle of a body of water or a vast desert, it should probably be discarded. Likewise, comparing the locations and attributes of other crimes will help narrow down the usefulness of the geographic profile.

12 Conclusions

The GRS method adapts the centrography method and Rossmo's formula. Unlike Rossmo's method, the GRS method is stable with respect to its internal parameters and unlike centrography, the GRS method is robust with respect to outliers. On the data sets we tested, the GRS method generated hot zones that contained the criminal's actual residence. Hence, we recommend the GRS method as a robust and stable model for creating a strong and effective model.

Appendices

A Stability Analysis Images

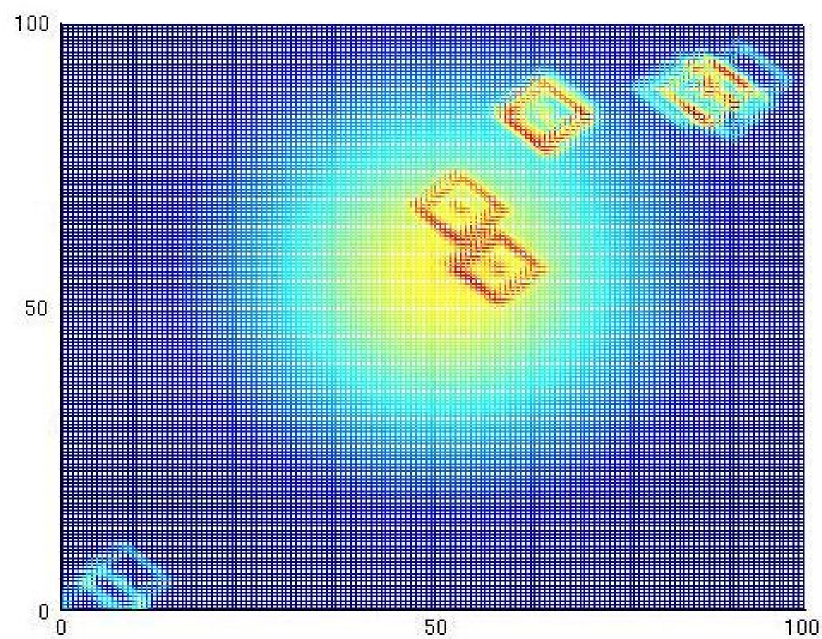

Figure 11: GRS Method on first eleven murders in the Yorkshire Ripper Case

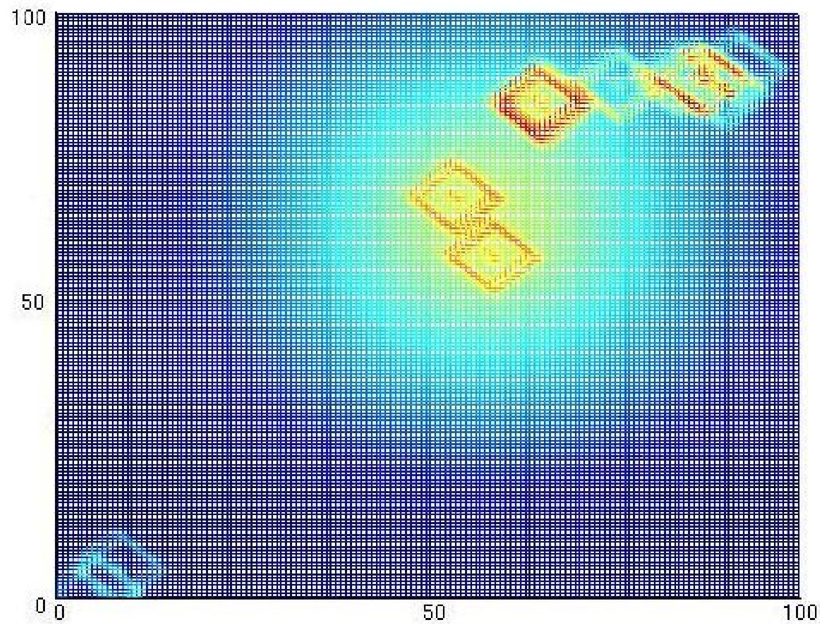

Figure 12: GRS Method on first twelve murders in the Yorkshire Ripper Case

References

[1] D. K. Rossmo, *Geographic profiling: Target Patterns of Serial Murderers.* PhD thesis, Simon Fraser University, 1995.

[2] B. Snook, D. Canter, and C. Bennell, "Predicting the home location of serial offenders: A preliminary comparison of the accuracy of human judges with a geographic profiling system," *Behavioral Sciences and the Law*, 2002.

[3] J. van der Kemp and P. van Koppen, "Fine-tuning geographical profiling," *Criminal Profiling: International Theory, Research, and Practice*, 2007.

[4] R. Kocsis and H. Irwin, "An analysis of spatial patterns in serial rape,arson, and burglary: The utility of circle theory of environmental range for psychological profiling," *Psychiatry, Psychology, and Law*, 1997.

[5] J. Warren, R. Reboussin, R. Hazelwood, and J. Wright, "The geographical and temporal sequencing of serial rape," *Journal of Interpersonal Violence*, 1991.

[6] L. W. Johnson and R. D. Riess, *Numerical Analysis.* Addison-Wesley Educational Publishers Inc, 1977.