

## HotSketch: Drawing Police Patrol Routes among Spatiotemporal Crime Hotspots

Alex Godwin  
Georgia Institute of Technology  
alex.godwin@gatech.edu

John Stasko  
Georgia Institute of Technology  
stasko@cc.gatech.edu

### Abstract

*During the course of a day, a police unit is expected to move throughout the city to provide a visible presence and respond quickly to emergencies. Planning this movement at the beginning of the shift can provide a helpful first step in ensuring that officers are present in areas of high crime, but these plans can quickly break down as they are pulled away to 911 calls. Once such an initial plan is deferred, police units need to be able to rapidly and fluidly decide where to go next depending on their immediate location and time. In this paper, we present our research to couple spatiotemporal analysis of historical crime data with sketch-based interaction methods. This research is presented through an initial prototype, HotSketch, which we describe through a set of use cases within the domain of police patrol route planning.*

### 1. Introduction


There are many approaches to proactive policing, but broadly, these practices are based on showing police presence, engaging with the community to learn their concerns, and analyzing historical crime reports to identify locations and people that are currently at increased risk of a crime. By moving through a police beat and creating a visible police presence, officers remind the public that they are nearby, discouraging potential violations while encouraging the lawful use of public space. By becoming involved with the community, an officer makes it easier for the residents to willingly participate in policing efforts. Finally, the use of crime analysis techniques and software can enable short-term tactical planning (e.g., which neighborhoods to patrol over the next hours and days) and long-term strategic efforts (e.g., drug market intervention programs that rehabilitate non-violent first term offenders).

While we are primarily interested in this last aspect for the research described in this paper, it is important to note that it facilitates the previous two. Commercial software for the analysis and prediction of criminal activity has seen steady deployment throughout the

country, bolstering and often improving on the capabilities of existing criminal analysts within departments. PredPol, for example, has been deployed in both Los Angeles and Atlanta to help officers determine at what location they should patrol [1]. Like many systems of this kind, it is not intended to support analysis of the path that an officer can take to get between destinations. Whether by car or on foot, hotspot analysis systems do not typically support the exploration of routes between locations.

While proactive policing is the intent, the reality is that many officers will spend a great portion of their day responding to 911 calls that take them away from their current location. Predictive crime systems can help reduce the response time for these calls by attempting to position officers in close proximity to areas that receive calls. This reactive aspect changes the nature of the plan, however, in that it puts the officer in a new location and causes time to elapse before a new plan can be created. It is no surprise that in many urban locations, the time of day can have a significant effect on the volume and types of crimes that are expected to occur. As the shift progresses and the unit responds to calls, the initial static analysis of the crime reports generated by many systems become stale. Police units could benefit greatly by a mobile system that allows them to view an updated analysis of crime hotspots based upon their changing location and time.

In this paper, we present our research to couple interactive spatiotemporal hotspot exploration with rapid route planning and analysis. Our primary components are: (1) a sketch-based approach to dynamic route planning, which allows a police officer to rapidly specify a path through the city without typing and review the volume and types of crime that occur along that route; and (2) a spatiotemporal hotspot approach that takes into account time of day, location, season, and recent event volume. Our primary contribution is the novel pairing of these components in an initial prototype, HotSketch, designed for mobile use on tablets. We present the details of this prototype and provide validation through a set of use cases within the



**Figure 1.** The HotSketch dashboard. A route has been drawn through a commercial shopping area on a Friday night at 11:30pm. The type of crimes along this route and at this time are presented in the category panel in the upper right. Below that, panels show the distribution for crimes by day of week, hour of day, and season (week of year) that are contributing to the count in the category panel. At the bottom of the dashboard, a histogram shows the distribution of the crimes along the planned route.

domain of police patrol route planning in the city of Atlanta.

HotSketch is designed to allow police officers to more rapidly utilize predictive models for crime in their neighborhood while away from the precinct, where expensive crime prediction software typically resides [2]. Rather than providing a static report, HotSketch is designed to allow officers to explore the data in their area through loose sketches that enable them to query the data and update the dashboard.


## 2. Related Work

There are many approaches to predictive policing, including massive data mining of text documents to identify the important named entities and establish potential criminal networks [3] [4]. In this paper, we are primarily interested in the spatiotemporal characteristics of crime, notably locations where a crime was reported by a member of the public or is known to have occurred. One of the most common analysis techniques for this form of event data is hotspot analysis, where historical records of criminal activity are used to predict areas where crime is likely to occur again due to the high frequency of past activity in that location. While there are a number of approaches for deriving hotspots [5], the most widely accepted and effective is Kernel Density

Estimation (KDE) [6], in which a function is used to smooth event locations across the map and aggregate areas of high frequency into easily perceived images.

This aggregation is controlled by a specified bandwidth, which specifies the extent to which an incident affects nearby areas. Selecting an appropriate bandwidth is challenging, and Maciejewski et al. have explored methods for creating variable spatial bandwidths in the analysis of crime data that help account for the difference in densely and sparsely populated regions [7]. For data that has a temporal component, such as crime, variations of the KDE approach can be applied to the temporal aspects of the data as well. Linked displays have been employed to capture the combination of spatial and temporal data by allowing the user to select identified hotspots on a map, which populates secondary linked views with representations of the temporal characteristics of the selected region (e.g., frequency for times of day) [8] [9]. By decomposing the individual elements of an event's time and date, Malik et al. were able to more accurately model the seasonal variation in critical events and make predictions for future occurrences [10].

KDE has also been employed in mobile applications specifically targeted at enhancing the capabilities of police officers on the move. Razip et al. designed a situation awareness app that provided officers with



**Figure 2.** Detailed view of crimes in the dashboard after zooming into the shopping center at the end of the path in Figure 1. Violent crimes are shown as red circles; non-violent crimes are rendered as smaller black circles. The heatmap is displayed behind the event circles to retain context. By tapping or clicking on the location of a crime, a detailed panel on the left reveals additional information. Below that, a street view is shown for the crime location.


mobile analysis capabilities for recent crimes in their area [11]. Mobile systems like this are typically easier to use for non-analysts, lower the technological barriers to more powerful forms of analysis, and reduce the need for additional technical staff at already cash-strapped public institutions [2].

One challenge with the tool described by Razip et al., and many tools like it, is that it does not provide the capability to make or analyze plans. For example, while the capabilities for exploring the characteristics of crime data in the immediate vicinity are present, what if an officer moves? At best, the officer can pan the map to a different location and explore the characteristics there. But what of the interstitial space, the path taken by the officer to get from the current location to the next? Visualizations abound for representing movement data through an area, though are primarily focused on determining patterns in the routes that people choose to take [12] [13] [14]. Instead, we are focused on the possible routes that an officer may take through a beat while on patrol. Turkay et al., provided an example approach for representing the characteristics of data that occurs along a line drawn across a map, though more for characterizing the changes in demographics across regions [15]. Drawing the line used for a query directly on the surface of the map, however, opens up opportunities for providing flexible interaction without

increasing the cognitive burden of learning a complex interface or querying language. By sketching the spatial queries rather than specifying them indirectly, it becomes much easier to form associations between both spatial and temporal features [16]. In a related domain, Forbus et al. demonstrated the utility for sketched maps in reasoning about a hostile battlespace by providing users with the capabilities to sketch complex tactical environments [17].

### 3. HotSketch

Our initial prototype is designed to run in the browser on tablets and laptops, and has been created using d3.js. The primary element is a 2D map centered on the current location of the user. Additional elements are available in hovering windows around the periphery of the map and are populated with data as the user interacts with the system. Figure 1 depicts the system interface. When the system is loaded, the officer is immediately provided with a heatmap of the crimes in the area based upon the current time of day and the geospatial location of the user. This heatmap consists only of the filtered events that occur within a relevant timeframe to the current time, which is determined through the approach described in Section 3.1. Then, the



**Figure 3.** An officer has drawn an initial route through her beat based on the driving directions given by a wayfinding app. There is a history of activity at the start and end of her route, but few incidents along the route. The distribution of violent crimes in this area are high, especially compared to the commercial shopping area in Figure 2. To provide police presence in high activity areas, she extends part of her trip through a busy area to the south.


officer can begin sketching potential routes through the area between any locations. As described in Section 3.2, these sketched paths can be used to query the filtered crimes to determine the types and locations of crimes that occur in proximity to the path. Finally, the officer can dig into the details of any particular crime by switching to the dot map view (Figure 2), which is described further in Section 3.3.

### 3.1 Hotspot Analysis

Our approach to hotspot analysis contends that there are observable patterns in the overall volume and types of crime that are dependent on when the analysis is being performed. For example, certain types of crimes occur more often at night than during the day, and the overall volume of crime increases during the hotter summer months. Similarly, the day of the week can influence this observable change in crime frequency, particularly between weekend and weekdays. Finally, as time progresses, the locations of hotspots will naturally migrate in response to the growth of the city and the effects of police intervening to alter this growth. To account for these factors, we use a weighted summation function that scores the temporal relevance of a crime to the current time and date when an officer opens

HotSketch. While a crime analyst would benefit from the ability to manually browse through range of potential parameters, we focus on providing functionality in the context of an operational environment in which the officer will want rapid answers rather than capabilities for model construction. The rationale is that if an officer opens the system on a Sunday morning in the middle of July, locations should be flagged as more relevant to the current geospatial crime landscape if they occurred in a similar temporal context. An officer opening the system on a Friday night in January should naturally see a different view.

To calculate the difference between the time a crime occurred and the present, we utilize a kernel function with a bandwidth parameter, in this case the bisquare function (Equation 1). For example, if a crime occurred on a Thursday and the current day is Friday, we want that crime to be considered less relevant than crimes that also occurred on Friday. Crimes occurring on more distant days of the week should be considered sparingly, if at all. We determine the difference between them as the number of days  $d_w$  and decide on a bandwidth size for which crimes can be considered relevant ( $\omega_w$ ). In most of the examples provided in this paper, we use ( $\omega_w = 2 \text{ days}$ ), and since Equation 1 increases monotonically as  $d_w$  decreases, this will heavily favor



**Figure 4.** While it does increase the length of her drive to her destination, the addition of a detour through a busy street has significantly increased her dwell time in areas of high activity. The trip histogram at the bottom now includes a new collection of activity for the detour. The distribution of violent crimes along this route are also high, particularly for aggravated assault and pedestrian robbery. While this distribution is similar to that of her original route, the volume of historical crimes along the new route is nearly double that of the former.

events that have occurred on the same week day while still including events that occur on the day before or the day after.

$$k(d, ) = \begin{cases} (1 - \frac{d^2}{2})^2, & d < \\ 0, & d \geq \end{cases}$$

**Equation 1.** Bisquare kernel

Continuing with this approach, we choose meaningful bandwidths for our other factors: time of day ( $d_t$ ), and difference in week of year, or season ( $d_s$ ). Like day of week, these measures are cyclical, so determining the difference parameter  $d$  must be done with care before using Equation 1. In most of the examples provided in this paper, we use ( $t = 3$  *ours*), as the day is divided into four overlapping shifts of six hours each. We choose ( $s = 6$  *weeks*), which favors events that occur in roughly the same season as the current date.


To determine the relevance that an incident has for the current time context, we create a weighted summation of the individual kernels. This function can be expressed as  $f(x_i, x_j)$  in Equation 2, in which  $x_i$  is the date of the incident and  $x_j$  is the current date. This process bears a resemblance to the Seasonal Trend

decomposition based on Loess (STL) approach utilized by Malik et al. [10], though there are some noticeable differences. In this example, we specify the individual weights to increase the importance of an event that occurs close to the same day of the week and time of day while deemphasizing the importance of the seasonal bandwidth. This weighting scheme favors our stated use case of an officer looking for highly contextual information based on current location and time. These parameters could be altered to reflect other use cases, however, such as determining routes of safe passage for school children throughout a season. Unlike Malik et al., we primarily rely on this summation approach to filter out irrelevant data from the heatmap and linked visualizations.

$$f(x_i, x_j) = \frac{2}{5}k(d_t, t) + \frac{2}{5}k(d_s, s) + \frac{1}{5}k(d_w, w)$$

**Equation 2.** Weighted summation of kernels

Once the equation has been applied to each event, the value for that event is added to a heatmap layer that is displayed on top of a 2D map (e.g., Figure 1). As the criminal incidents are not counted equally, this leads to variation in the distribution of events from a standard heatmap and movement of hotspot locations from a



**Figure 5.** Panels provide a detailed look at the distribution of relevant events through histograms. For each panel, the current time is filled in black; other bars are colored by the score for individual kernel windows (darker bars score higher than light bars).

straightforward count. Instead, this new distribution reflects both the location of events and the temporal context in which the request was made, allowing an officer to see the distribution of events for their area and throughout their shift.

### 3.2 Sketching Patrol Paths


Paths are drawn onto the map using editing controls in the top left of the dashboard. Using a mouse or by tapping on the map, an officer creates a path by specifying the control points that define a polyline. Once the path has been specified, the user can choose to edit the path by dragging the control points or to delete it entirely. Once finalized, a completed path can be explored in the context of the data by clicking on it. HotSketch will then compare the line segments of the path to the crime incidents in the area to determine which ones are closest. In addition, we derive the point on the line that is closest to each point that is within a minimum distance specified by the user. In the examples presented in this paper, this threshold is set to 200m, which is sufficient to include incidents within one or two blocks of the path drawn by the user.

By determining the point along the line that is closest to each crime, we can reconstruct the order in which an officer would pass by these crimes from the beginning of the patrol path to the end. In Figure 1, for example, the officer has drawn an initial route from a location into a commercial district with numerous shopping options. Once the details of the route have been loaded into the dashboard, the officer can review the types of crimes that occur along the route. The events are filtered and weighted based on the approach described in the previous section, emphasizing those crimes that occur not only in proximity to the line sketched by the officer but in relevance to the temporal context in which the path is created.

That is, the time and date in which the user is drawing these lines affects the types of data shown (subject to the approach described in the previous section). For example, in Figure 1 the panel in the upper right indicates the relative distribution of crimes along this route by category. This panel depicts only the crimes that have a non-zero value from Equation 2. The panels below indicate the frequency of events according to day of week, time of day, and week of year. Since Equation 2 is additive, events that receive a low or zero score for any of the individual kernel components can still appear as long as they have a positive value for one of the other kernels. The emphasis of the kernels is evident, however, in the peaks of the distributions for each of the detailed panels. We also shade the fill color of the bars within the individual histograms to indicate the current day, hour, and week (Figure 5). The darkest bar indicates the current time, and a monotonically decreasing grayscale indicates the potential score of events coinciding with the other bars for each kernel.

### 3.3 Event Details

If the officer wants to review additional details, she can switch to a dot map that shows the exact location of all the crime events within the area that have not been filtered out by the approach described in Section 3.1. In Figure 2, for example, an officer is examining relevant crimes that have occurred in the proximity of a commercial shopping area. In the figure, the most common type of incident is a larceny from an automobile, which likely results from the massive parking lot within the facility and the propensity for shoppers to leave valuables visible and unattended in their cars. When the officer clicks on one of the event circles on the map, linked displays to the left provide details for the selected event. The top pane provides text details on the event as they are available, such as the type of crime, the date, the time of day, and a description if one is available. To help provide additional visual context, a secondary pane on the map provides street-



**Figure 6.** The officer switches views to get a detailed look at a neighborhood with a history of frequent criminal activity. By clicking on the dots representing the relevant crimes, the officer becomes familiar with the details of past incidents and the context of the locations in which they occurred.

level imagery pointed at the crime location. This imagery is obtained through the Google Street View Image API, and reflects the most up-to-date picture available at that location regardless of the date or time the crime occurred.

**Table 1.** Distribution of Crimes in Atlanta

<b>Violent</b>	
Aggravated Assault	17,500
Robbery-Pedestrian	13,364
Robbery-Residential	1,739
Robbery-Commercial	1,697
Rape	841
Homicide	519
<b>Non-violent</b>	
Larceny-From Vehicle	69,611
Larceny-Non Vehicle	59,611
Burglary-Residence	40,400
Auto Theft	35,325
Burglary-Nonresidential	7,778

## 4. Evaluation

To validate our approach for providing officers with the capability to rapidly utilize predictive models for crime in their neighborhood while away from the precinct, we loaded representative data from a major

metropolitan area into HotSketch. Then, to determine the system’s ability to facilitate analysis and comparison of alternative routes through space, we explored a usage scenario in which an officer is moving through the city. While we envision eventually performing more robust validation of our system in a representative deployment context to establish ecological validity, the approach presented here provides promising, albeit preliminary, indication of the system’s usefulness.

### 4.1 Data

The data used in the evaluation of HotSketch is derived from publicly available crime data provided by the city of Atlanta. This data, known as Uniform Crime Reporting (UCR), consists of the major types of crime that the city reports to the FBI on a regular basis. From the initial UCR data, we analyzed nearly 250K incidents from 2009 to mid-2016. The incidents can be violent crimes against people or non-violent crimes. The distribution of criminal activity during this period is given in Table 1.

### 4.2 Usage Scenario

As evidenced in Table 1, non-violent property crimes such as larceny from a vehicle account for a

massive portion of the crimes throughout the city. While the police cannot be everywhere at once, maintaining a visible presence in areas of significant property crime can reduce the number of incidents at that location.

In our use case, an officer is moving through her beat as part of her normal patrol. Using a wayfinding app, she is given a relatively quick route from her current location to her destination within the eastern edge of her beat. However, by quickly drawing the route on HotSketch, she sees that there is an area directly to the south of her that appears to have a high volume of activity (Figure 3). The area contains a collection of single family homes and small apartment buildings with a history of assault, robbery, and larceny. While this new location is a little out of the way for her, it is also an opportunity to establish a presence and spend some time in an area within her patrol beat without being summoned there for a call. She quickly sketches a few alternatives, deciding on a route that detours through a section of the neighborhood before returning to her original route (Figure 4). By taking this detour, however, she is substantially increasing her dwell time in proximity to areas of historically high crime, and the timeline panel along the bottom of the dashboard reflects an increased gain in nearby incidents towards the middle of her trip.

To familiarize herself with the problem areas along her new route, the officer switches to the details view in the dashboard (Figure 6). In this image, she has zoomed into a residential area where a high volume of relevant crimes have occurred. Clicking on one of the violent crimes provides her with a description – an aggravated assault at the address, one of many in the vicinity. The street view panel has imagery, so she is able to see the building in which many of these incidents are occurring. She now has a route, and knows what to look for on her way through. She will keep an eye on this location in the future, and look for opportunities to get to know residents in the area so that her presence as an officer of the law will be recognized and respected.

## 5. Conclusions and Future Work

In this paper, we described our preliminary research to allow police officers to more rapidly utilize predictive models for crime in their neighborhood while away from the precinct. We described the main aspects of our approach, which couples interactive spatiotemporal hotspot exploration with rapid route planning and analysis.

There are several promising directions for future research in this area. First, while our initial validation provides evidence that HotSketch has the capability to facilitate exploration within the confines of a narrowly

defined task, it would be beneficial to pursue a more thorough evaluation by putting the prototype directly into the hands of police officers. This would allow us to not only evaluate the efficacy of the current set of capabilities provided by the prototype, but also to elicit additional potential design requirements from officers directly.

Additionally, it would be useful to design a set of capabilities within HotSketch that are directed towards the needs of civilians. For example, a publicly available version of the prototype could be adapted to allow citizens to explore crime and other types of data within their community. The sketching capabilities provide an excellent foundation for non-expert interaction, and could be adapted to a host of new capabilities for allowing citizens to query and explore spatiotemporal data.

## 6. Acknowledgements

This work was supported in part by the Department of Homeland Security's VACCINE Center of Excellence. This work was also supported in part by the National Science Foundation. The authors wish to thank the Atlanta Police Department for providing updated and open access to the crime data.

## 7. References

- [1] "PredPol: The Predictive Policing Company," [Online]. Available: <http://www.predpol.com/>.
- [2] R. E. Roth, K. S. Ross, B. G. Finch, W. Luo and A. M. MacEachren, "A user-centered approach for designing and developing spatiotemporal crime analysis tools," in *Proceedings of the Sixth International Conference on Geographic Information Science (GIScience '10)*, Zurich, Switzerland, 2010.
- [3] H. Chen, W. Chung, J. Jie Xu, G. Wang, Y. Qin and M. Chau, "Crime data mining: a general framework and some examples," *Computer*, vol. 37, no. 4, pp. 50-56, 2004.
- [4] H. Chen, D. Zeng, H. Atabakhsh, W. Wyzga and J. Schroeder, "COPLINK: managing law enforcement data and knowledge," *Communications of the ACM*, vol. 46, no. 1, pp. 28-34, 2003.
- [5] J. Eck, S. Chainey, J. Cameron and R. Wilson, "Mapping Crime: Understanding hotspots," U.S. Department of Justice, Washington, DC, 2005.
- [6] S. Chainey, L. Tompson and S. Uhlig, "The utility of hotspot mapping for predicting spatial patterns of crime," *Security Journal*, vol. 21, no. 1, pp. 4-28, 2008.
- [7] R. Maciejewski, S. Rudolph, R. Hafen, A. Abusalah, M. Yakout, M. Ouzzani, W. S. Cleveland, S. J. Grannis and D. S. Ebert, "A visual analytics approach to understanding spatiotemporal hotspots," *IEEE*

- Transactions on Visualization and Computer Graphics*, vol. 16, no. 2, pp. 205-220, 2010.
- [8] S. Afzal, R. Maciejewski and D. S. Ebert, "Visual analytics decision support environment for epidemic modeling and response evaluation," *Visual Analytics Science and Technology (VAST)*, pp. 191-200, 2011.
- [9] C. Brunson, J. Corcoran and G. Higgs, "Visualising Space and Time in Crime Patterns: A comparison of methods," *Computers, Environment, and Urban Systems*, vol. 31, no. 1, pp. 52-75, 2007.
- [10] A. Malik, R. Maciejewski, S. Towers, S. McCullough and D. S. Ebert, "Proactive Spatiotemporal Resource Allocation and Predictive Visual Analytics for Community Policing and Law Enforcement," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 20, no. 12, pp. 1863-1872, 2014.
- [11] A. M. Razip, A. Malik, S. Afzal, M. Potrawski, R. Maciejewski, Y. Jang, N. Elmqvist and D. S. Ebert, "A Mobile Visual Analytics Approach for Law Enforcement Situation Awareness," in *PacificVis*, Yokohama, Japan, 2014.
- [12] G. Andrienko, N. Andrienko, J. Dykes, S. I. Fabrikant and M. Wachowicz, "Geovisualization of dynamics, movement and change: key issues and developing approaches in visualization research," *Information Visualization*, vol. 7, no. 173, pp. 3-4, 2008.
- [13] N. Andrienko and G. Andrienko, "Spatial generalization and aggregation of massive movement data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 2, pp. 205-219, 2011.
- [14] C. Tominski, H. Schumann, G. Andrienko and N. Andrienko, "Stacking-based visualization of trajectory attribute data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2565-2574, 2012.
- [15] C. Turkay, A. Slingsby, H. Hauser, J. Wood and J. Dykes, "J. (2014). Attribute Signatures: Dynamic Visual Summaries for Analyzing Multivariate Geographical Data," *IEEE Transactions on Visualization and Computer Graphics*, 2012.
- [16] A. D. Blaser, *Sketching Spatial Queries*, Munster, Germany: Dissertation, University of Munster, 2000.
- [17] K. D. Forbus, J. Usher and V. Chapman, "Qualitative spatial reasoning about sketch maps," *AI Magazine*, vol. 25, no. 3, p. 61, 2004.