HunchLab Demonstration Project for the New York City Police Department

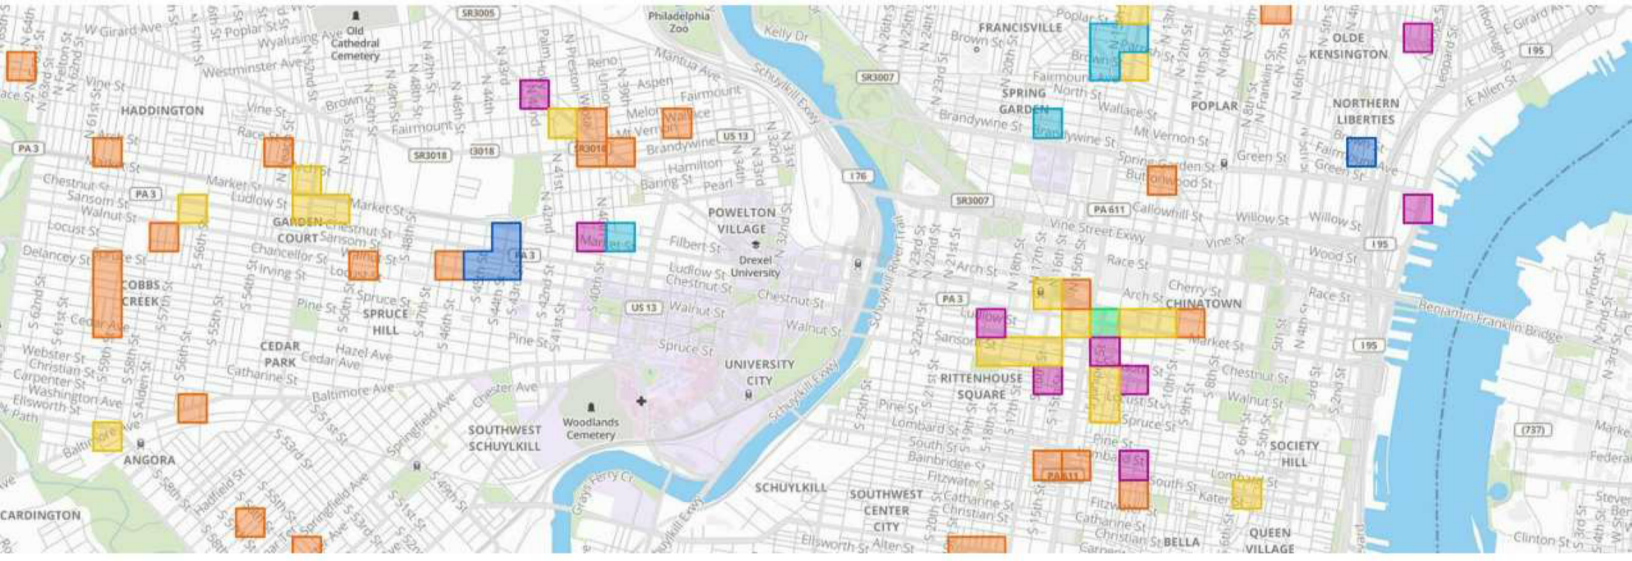

HunchLab

Prepared by

azavea
340 N 12th St, Suite 402
Philadelphia, PA 19107
(215) 925-2600
http://www.azavea.com

Table of Contents

Executive Summary

Azavea's HunchLab system provides next generation predictive policing technologies in an easy-to-use web-based solution. Built to utilize a broad set of crime theories and datasets, HunchLab automatically generates robust forecasts for crime and distills these forecasts into target mission areas for easy deployment in the field. We are offering the New York City Police Department the opportunity to pilot HunchLab in an operational setting for 24 months, which will enable the Department to determine HunchLab's efficacy as a daily tool used by the NYPD.

HunchLab is delivered as a secure cloud subscription, which reduces operational cost and complexity. Azavea will manage the hosting infrastructure, security updates, and 2nd tier support. This cloud-based approach enables the HunchLab to leverage significant amounts of computing power in an elastic manner, a critical requirement for providing the advanced statistical algorithms the system employs. To provide a secure application, we have consulted the FBI's CJIS guidelines within our system architecture, including mitigation techniques such as encrypting data in transit and providing optional 2-factor authentication. More information about HunchLab security is provided in the Appendix at the end of this proposal.

Current trends reflect an interest in predictive policing, increasing adoption of cloud services, pervasive location information, and the proliferation of mobile devices. These trends are occurring in an extremely constrained budget environment in most law enforcement agencies. Azavea believes these factors will combine to cause the law enforcement community to undergo a transformational change in the use of technology, similar in scope to the advent of GIS-based crime mapping in the 1990's. The new HunchLab anticipates these advances in an unprecedented manner; we believe that HunchLab 2.0 truly represents the future of policing. We look forward to working with the New York City Police Department to implement our current technology and spearhead new, innovative approaches to improve public safety.

Demonstration Project

Azavea believes that, where possible, vendors should provide low-cost opportunities for governments to validate solutions before committing to long-term acquisition. To support this belief, we have designed a custom demonstration project for NYPD that provides experience with HunchLab and validation of adoption within the agency. The demonstration project begins with a short configuration period in which settings are tweaked and the system is aligned with the department's priorities. The demonstration project provides 24 months of subscription. During this period, the New York City Police Department can use HunchLab operationally within up to three police precincts. Predictions will also be available for some areas outside of these three operational precincts for purposes of accuracy evaluations.

Phase 1: Configuration and Vetting

The initial configuration will commence once we receive historic crime data. Over the course of the configuration timeframe, Azavea will work with NYPD to make sure the mission areas make sense, to tweak settings as necessary so the system reflects the priorities of the department, and to measure the initial accuracy of the system.

Phase 2: Operational Use

The operational use will begin immediately after the configuration and vetting phase. During this phase, HunchLab will be deployed in up to three police precincts of the Department's choosing. Command staff will task officers to focus on mission areas. Mission areas will be disseminated as deemed most appropriate by NYPD (through the web application, electronic PDFs, or physical print outs).

Optional Activities

There are several additional activities NYPD may want to conduct during the demonstration project. The HunchLab API will be fully accessible, allowing your IT department to pull data out of HunchLab and into other mapping or analysis applications such as the Domain Awareness System. You might also want to collaborate with an academic during this period to examine the adoption, impact, and accuracy of the tool.

The team at Azavea that works on HunchLab also looks forward to the feedback that NYPD can provide about the application. This feedback may take the form of user surveys, conference calls, visits to the precinct, and officer "ride-alongs" as deemed appropriate by NYPD. Our past pilot implementations of HunchLab have been a rich source of enhancement ideas and we have often delivered such enhancements before the pilots have concluded.

Project Deliverables

The demonstration project will deliver a view into the operational impact of predictive analysis at your department. It will also validate the technical requirements of such systems in regards to data formats, data cleanliness, and timely dissemination of information to the field. The HunchLab pilot program provides insight into the expected level of accuracy for predictive modeling of your data through automated accuracy measurements from the HunchLab modeling process. Finally, summary information about the factors represented in the HunchLab models for each crime type will be provided.

Project Features

- Automated advanced crime forecasts using crime data and other related data sets
- Interactive modern desktop browser interface to configure and view mission areas
- On-demand generation of PDF reports for mission areas for print and email distribution
- Adaptive mobile browser interface to HunchLab (as available based upon NYPD roll out of devices and tweaks by Azavea)
- User single sign-on via the SAML standard
- Access to HunchLab Advisor as it become available

Project Assumptions

- HunchLab will be provided for a period not to exceed 24 months (configuration and operational periods) unless extended by the mutual agreement of both parties.
- HunchLab will be utilized operationally in up to three police precincts.
- HunchLab will be hosted in Azavea's secure Amazon Web Services cloud infrastructure with data stored only in data centers on United States soil.
- NYPD will provide access to key crime data fields in an open format for a five-year period – either in a specified CSV or in a similar, accessible format that Azavea can transform into such CSV such as an ODBC connection.
- Daily (or more frequent) crime data uploads to HunchLab will be posted out of NYPD's internal network automatically via a scheduled task or similar process. Azavea can support the development of this process or NYPD's IT department can handle the process.
- NYPD users will access the application from modern web browsers that support secure connections as specified in the application requirements provided below.
- Azavea will not utilize NYPD data for any purpose other than to provide this project.
- Azavea will purge NYPD data from HunchLab upon NYPD request at any time.

Demonstration Project Cost

Item	Total
Setup & Integration Fee	[redacted]
1st Year Annual Service & Connection Fees (First 6 Months Free)	
2nd Year Annual Service & Connection Fees	
Total	

Our quoted price is $[redacted] for this demonstration project in up to three New York precincts. Out of consideration for the nature of this demonstration and the feedback we will receive through this project, this quoted price represents a discount of [redacted] over our standard retail pricing of [redacted]

In addition to the above discount, there are additional aspects of this project that would typically incur additional services cost. We are not charging for this additional work. While the demonstration project is limited geographically, the setup work will likely remain at the same order of magnitude as a full

implementation. NYPD is a sophisticated client that will likely exercise HunchLab in ways not applicable to typical jurisdictions. NYPD has expressed the desire to leverage many distinct data sets including complaints, calls for service, arrests, GPS feeds, and geographic layers. We expect to spend significant time incorporating such data sets into the demonstration project along with the relevant security requirements of the department. NYPD has also expressed the desire to model crime at a smaller geographic unit than the default 250m resolution. While HunchLab currently supports a customizable resolution, we have not explored the limits of the modeling system in this regard. We intend to do so in order to support unique environments such as New York City. NYPD also presents unique challenges related to the verticality of the city. We will likely be building more statistical models than we normally would in order to separate ground-level activity from activities both above and below ground. NYPD has also expressed the desire to consume our predictive layers within existing systems such as the Domain Awareness System, which will warrant some support on Azavea's part.

Milestone Payments

Payments for HunchLab will be tied to project milestones, outlined below:

- Setup Milestone (Setup Fee, Connection Fees, and 1st Annual Service Fee)
    - This milestone is complete once Azavea has completed the initial setup and NYPD has vetted the missions being generated by HunchLab. Azavea will provide up to 6 weeks for vetting the mission areas, crime models, and resource configuration settings once the mission areas are available for consumption. After the mission areas have been recognized as vetted by the client or have been available for 6 weeks – whichever occurs first - the system will be considered vetted and an invoice will be issued. In other words, the annual service fee period will begin on the first day that missions were provided to the client (which indicates that base configuration is complete) but the client will not be invoiced until the vetting process has been completed.
- Additional Yearly Subscription Milestones (Connection Fees, Annual Service Fee):
    - Subscription fees for subsequent years will be invoiced at beginning of each additional year of service.

Implementation Timeline

Azavea has delivered a configured HunchLab system from start to finish in under a week. Most deployments, however, benefit from additional time for the client to vet the configuration decisions being made within the system, to validate the resulting missions, and to educate staff as the client progresses toward operational use of the tool. A more typical timeline lasts 2-3 months. Below is a sample work plan that was designed in collaboration with another client:

Sample Project Plan

This work plan is broken down into two-week sprints to reflect an agile SCRUM methodology. For instance, Sprint 1 would correspond to work done during weeks 1 and 2.

When	What	Who	On/Off Site
Sprint 1	**OBJECTIVE: Project Planning & Integration Requirement Fulfillment**		
	Kickoff meeting	**Both**	**Both**
	Crime data requirements meeting	**Both**	**Both**
	Contextual geographic/temporal data requirements meeting	**Both**	**Both**
	Authentication requirements meeting	**Both**	**Both**
	Define CSV template or database view structure	**Both**	**Both**
	Create crime database view or CSV export routine	**Client**	**On**
Sprint 2	**OBJECTIVE: Integration configuration**		
	Create integration utility	**Azavea**	**Off**
	Create HunchLab administrative accounts for client	**Azavea**	**Off**
Sprint 3	**OBJECTIVE: Live Data Integration**		
	Deployment of integration utility	**Both**	**On**
	Integration utility support	**Azavea**	**Off**
	Validation of imported event data	**Both**	**Both**
	Validation of authentication system	**Both**	**Both**
	Administrative Training (1st Pass)	**Azavea**	**Off**
	Configure system-level settings	**Azavea**	**Both**
Sprint 4	**OBJECTIVE: Model Configuration**		
	Configure crime classifications	**Both**	**Both**
	Configure polygon hierarchies	**Both**	**Both**
	Configure models	**Both**	**Both**
	Configure police resources	**Both**	**On**
	Configure shifts	**Both**	**On**
	Configure contextual geographic variables (optional)	**Azavea**	**Both**
	Support configuration work	**Azavea**	**Off**

Sprint 5	**OBJECTIVE: System training & validation**		
	Administrative / Analyst Training (2nd Pass)	**Both**	**Both**
	Command Training (1st Pass)	**Both**	**Both**
	System validation	**Client**	**On**
	System remediation	**Azavea**	**Off**
Sprint 6	**OBJECTIVE: Go-live**		
	End-user Training (1st Pass)	**Azavea**	**Off**
	Record Video of End-User Training (1st Pass)	**Azavea**	**Off**
	Command Training (2nd Pass)	**Azavea**	**Off**
	System signoff	**Client**	**On**
Sprint 7+	**OBJECTIVE: Post-Deployment**		
	Configure contextual temporal variables (optional)	**Azavea**	**On**
	Mobile interface testing (optional)	**Both**	**On**
	User feedback sessions	**Both**	**On**

Support & Service Level Agreement (SLA)

Service Level

During the Term of the HunchLab subscription, Azavea shall make commercially reasonable efforts to maintain the operation and availability of the application to the Customer at least 99.9% of the time as measured over the course of a calendar month. This SLA level corresponds with approximately 44 minutes of unplanned downtime per month. Scheduled downtime will not be included in these calculations but generally will amount to less than 1 hour per month. If the application does not meet the SLA, the Customer may request Service Credit as described below. This SLA states the Customer's sole and exclusive remedy for any failure by Azavea to provide the service.

Definitions

The following definitions shall apply to the SLA:

- *"Downtime"* means that the application and API are unavailable as measured by availability and a valid response being received from an external monitoring service maintained by Azavea.
- *"Downtime Period"* means a period of two consecutive minutes of Downtime. Intermittent Downtime for a period of less than two minutes will not be counted towards any Downtime Periods.
- *"Scheduled Downtime"* means those times where Azavea notifies Customer of periods of Downtime at least five calendar days prior to the commencement of such Downtime. There will be no more than eighteen hours of Scheduled Downtime per calendar year. Scheduled Downtime is not considered Downtime for purposes of this SLA, and will not be counted towards any Downtime Periods.
- *"Monthly Uptime Percentage"* means total number of minutes in a calendar month minus the number of minutes of Downtime suffered from all Downtime Periods in a calendar month, divided by the total number of minutes in a calendar month. For purposes of the Server Uptime Level, a lapse in server availability is calculated from the time Azavea detects or otherwise becomes aware of an incidence of a service interruption and ending when the service is restored, regardless of where the outage originated.
- *"Service Credit"* means the following:

Monthly Uptime Percentage	Service Credit added to the end of the Service Term, at no charge to Customer
< 99.9% & ≥ 95.0%	1 additional week added to the subscription
< 95.0%	2 additional weeks added to the subscription

In order to receive any of the Service Credits described above, Customer must notify Azavea within thirty days from the time Customer becomes eligible to receive a Service Credit. Failure to comply with this requirement will forfeit Customer's right to receive a Service Credit.

Service Credits

Service Credits may not be exchanged for, or converted to, monetary amounts.

Monitoring

Azavea shall maintain a monitoring service external to the data center housing the equipment supporting the application and shall monitor the service with reasonable frequency and duration.

Scheduled Downtime

Azavea shall provide at least 5 calendar days or more notice to customers if there is planned downtime. Tasks performed during planned downtime may include:

- Application of security patches
- Upgrades to software
- Updates to the database
- Other activities as necessary to maintain the integrity, stability and performance of the web services.

SLA Exclusions

The SLA does not apply to unavailability or any performance issues:

(i) caused by factors outside of Azavea's reasonable control including without limitation, acts of God, acts of government, flood, fire, earthquakes, civil unrest, acts of terror, strikes or other labor problems.; or

(ii) caused by a malicious internet attack including a denial of service attack; or

(iii) that resulted from Customer's equipment or activity or third party equipment or activity, or both (not within the primary control of Azavea)

Changes to Service

Azavea may make commercially reasonable modifications to the Service, or particular components of the Service, from time to time. Azavea will use commercially reasonable efforts to notify Customer of such changes.

Customer Support

Azavea shall use commercially reasonable efforts to provide the following support services for customers:

- Provide telephone, web and/or email support to customers during normal business hours (9am – 6pm, Eastern Time); and
- Respond to customer support queries regarding within one to five business days, depending on the severity of the issue.

HunchLab Application Features

HunchLab 2.0 incorporates concepts such as: temporal patterns (time of day, day of week, day of month, seasonality); weather; risk terrain modeling (locations of bars, bus stops, etc.); socioeconomic indicators (such as collective efficacy indicators); historic crime levels; and near-repeat patterns to help police departments understand and respond more effectively to crime using the resources available to them. After predicting crime expectations across the entire jurisdiction for a shift, HunchLab calculates the relative crime level per unit of patrol effort for each area, sorts these areas from highest to lowest relative risk, and selects the mission areas that can be patrolled by the available resources. This process maximizes the impact of patrols – placing them in areas where crime risk is greatest – and ensures that the right quantity of mission areas is crafted.

At the core of HunchLab 2.0 is a new crime-forecasting engine. These forecasts power the predictive missions feature, enabling departments to proactively generate the appropriate quantity of mission areas based upon the organizational and societal importance of various types of crime. The forecasting engine uses ensemble machine learning approaches that can incorporate the following crime patterns into a single prediction of criminal risk:

- Baseline crime levels
  - Similar to traditional hotspot maps
- Near repeat patterns
  - Event recency (contagion)
- Risk Terrain Modeling
  - Proximity and density of geographic features (points, lines, and polygons)
- Routine activity theory
  - Offender: proximity and concentration of known offenders
  - Guardianship: police presence (historic AVL / GPS data) [Requires additional integration work not available in a pilot.]
  - Targets: measures of exposure such as population, parcels, or automobiles
- Collective Efficacy
  - Socioeconomic indicators, neighborhood heterogeneity, etc.
- Temporal cycles
  - Seasonality, time of month, day of week, time of day, etc.
- Recurring temporal events
  - Holidays, sporting events, etc.
- Weather
  - Temperature, precipitation, etc.

Azavea believes that the use of non-crime data sets as variables within an operational crime prediction system is important because variables based solely upon crime data become skewed as predictions are used operationally. As crimes are prevented in mission areas due to police response, the only variables identifying areas as high risk are skewed in other systems. By including other data sets, our system is more robust against this issue.

While available GIS tools have enabled law enforcement agencies to advance their understanding of crime through more effective geographic visualization for many years, these tools have traditionally required trained analysts, are used by a tiny subset of agency personnel, and are largely reactive in nature. There is a compelling need for new crime analysis tools that perform automated discovery of crime patterns and make that information available in a format that can be easily understood and acted upon at a variety of agency levels, including officers in the field. HunchLab provides these capabilities. It is a groundbreaking system that "learns" which crime theories matter for a given location and automatically calibrates the influence of these theories – both individually and as arbitrarily complex interactions – within a designated geography to identify and address a wide range of crime and other public safety issues.

1. Predictive Missions

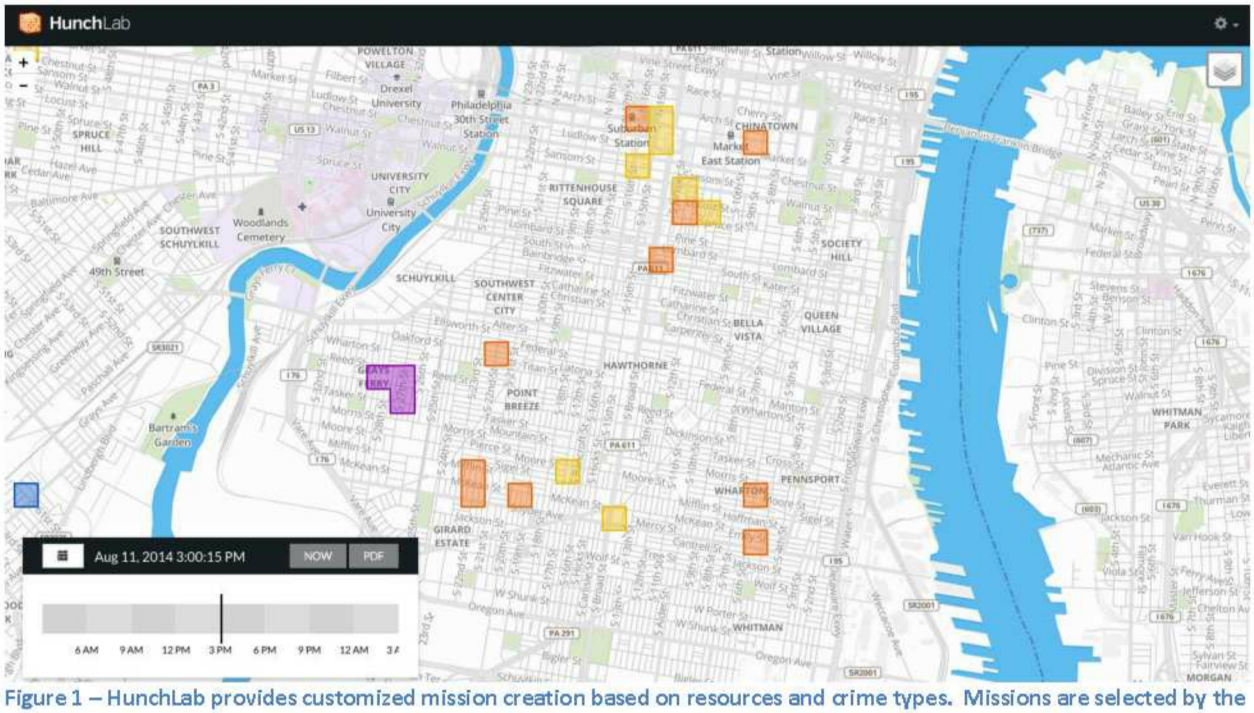

Figure 1 – HunchLab provides customized mission creation based on resources and crime types. Missions are selected by the combined, weighted risk of all configured crime models. Color represents the dominant risk for a mission area.

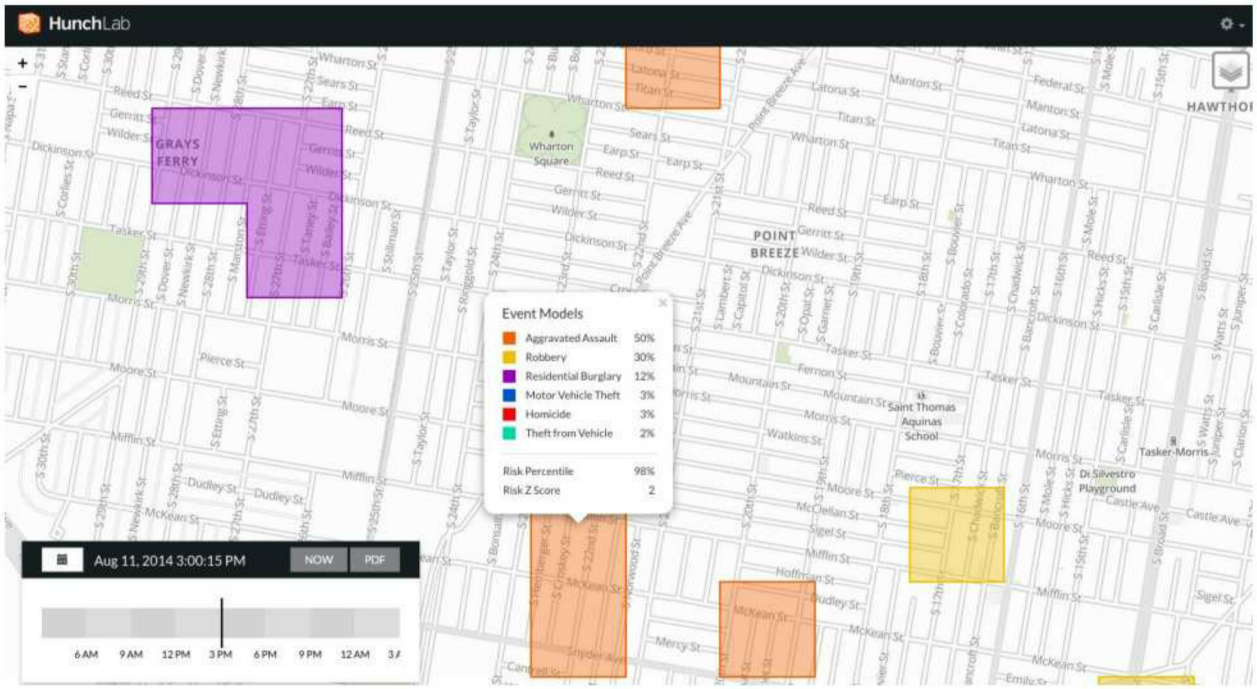

Figure 2- Each mission area displays a risk profile of the crime types that went into selecting this location.

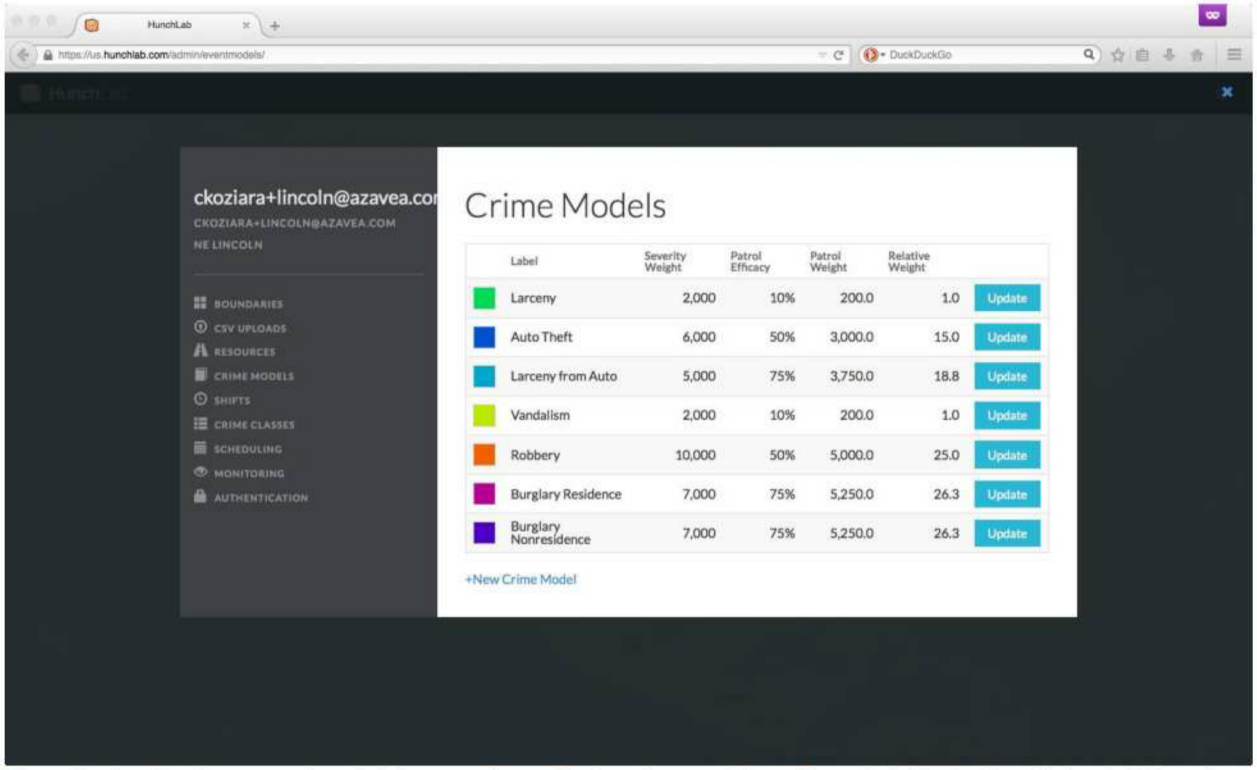

Figure 3 – A police department's priorities are reflected in the crime models configured within the HunchLab administrative user interface. Severity weights enable the department to tell HunchLab how important it is to prevent each type of crime. In this example the cost of crime numbers from the RAND Corporation are utilized to align policing priorities to the societal impact of crime. Patrol efficacy values enable the department to specify how much impact they believe patrols will make on each type of crime. The result is that missions show up where the most important, preventable crimes are likely to occur.

2. Sample Mission PDFs

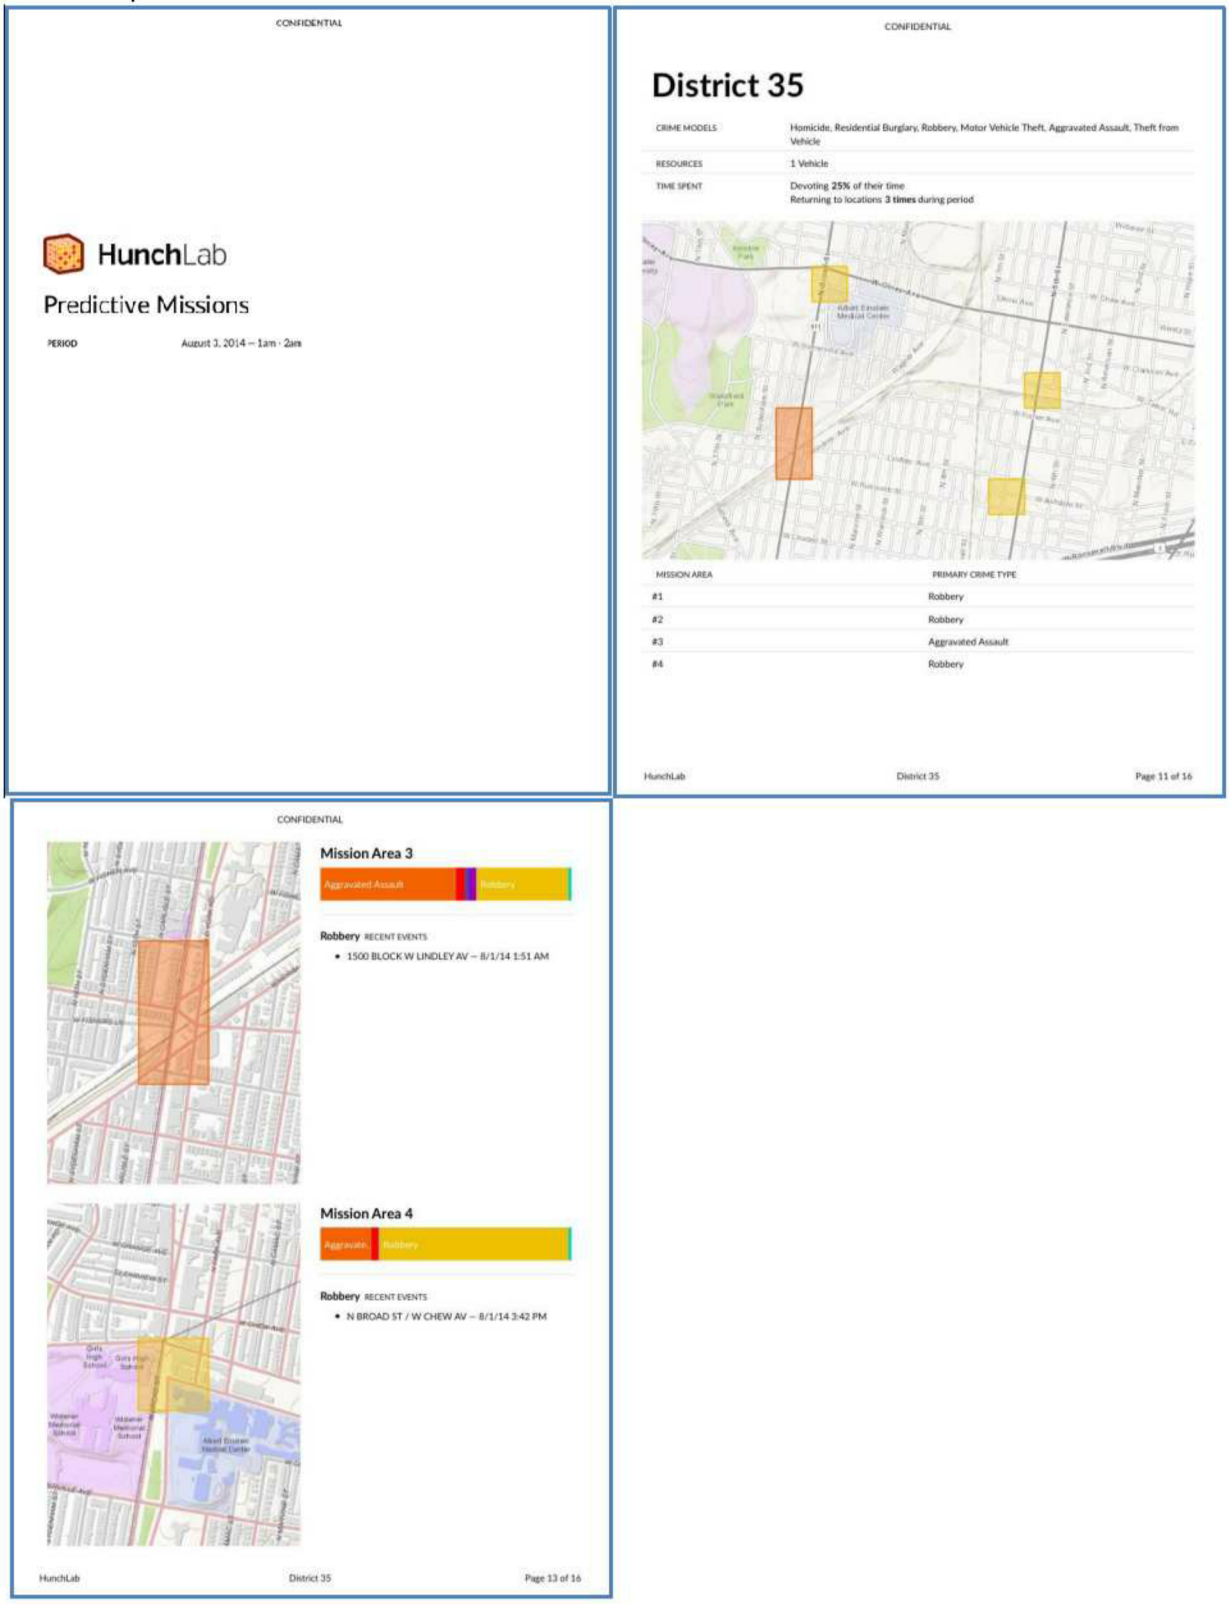

Figure 4 – Mission data can be printed or distributed as PDFs and taken on patrol if network connectivity is not available.

3. Example Model Composition

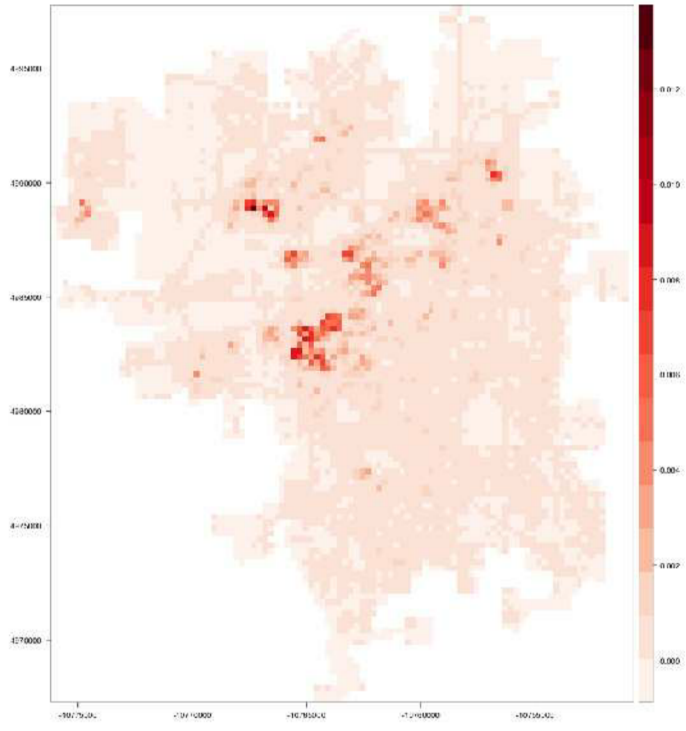

Figure 5 – HunchLab forecasts the expected count for each configured crime type within a shift for each small location within the jurisdiction. This figure shows an example of the map of one such set of forecasts.

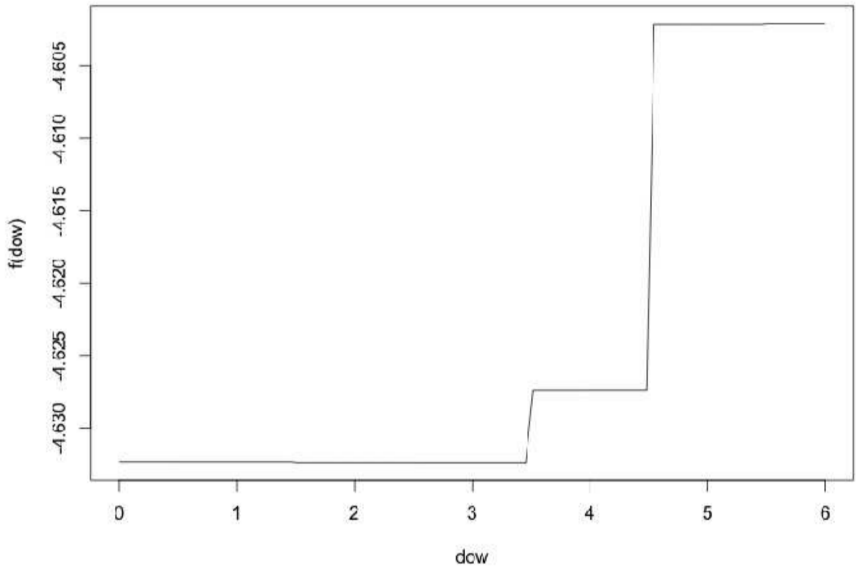

Figure 6 – The forecasting models can be examined to visualize what the system has determined effects the risk levels. In this case, the system learned how Friday, Saturday, and Sunday (4, 5, 6) have higher levels of assaults in Philadelphia.

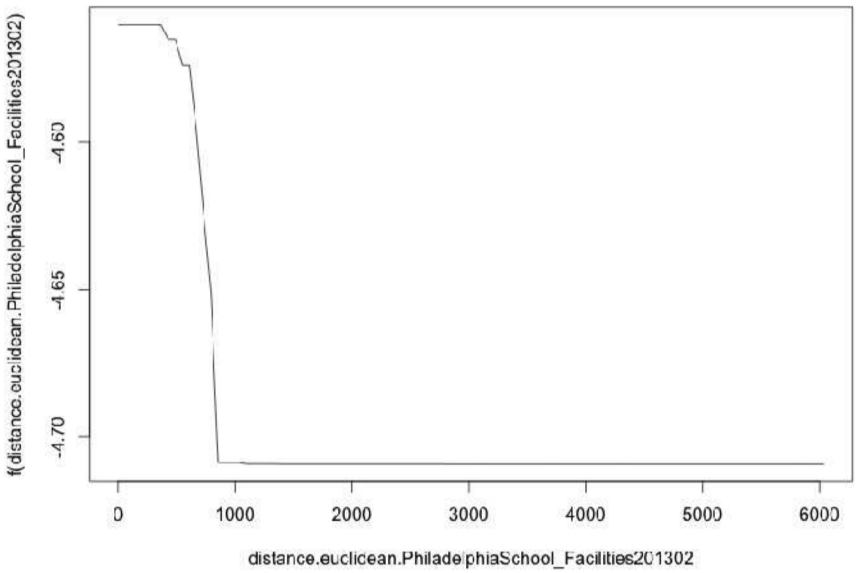

Figure 7 – In this example, the proximity to schools is shown to increase assault risks extending to about 900 meters.

Figure 9 – In this example, the proximity to schools is shown to increase assault risks extending to about 900 meters.

Figure 10 – A visualization of the different data sources contributing to forecasts for Motor Vehicle Theft. HunchLab's forecast in each city is composed of the same data types, but uses them in each individual model to varying degrees based upon the local data.

Related Activities: HunchLab Advisor

HunchLab Advisor is a software service that enables departments to engage in evidence-based policing through the Field Test, Experiment, and Adaptive Tactic initiatives of the module. This module is currently under development by Azavea with statistical design completed and user interfaces being developed. As the module becomes available, Azavea will provide access to this functionality.

Field Test

The Field Test component shows departments the impact of new initiatives within a specified geographic area. For example, if a department adopts a new DDACTS deployment model in a particular part of the jurisdiction, the Field Test can tell the department the likely impact of the new program. This output includes information like how many crimes the new strategy likely prevented.

Experiment

The Experiment component gives departments an easy way to set up a Randomly Controlled Trial (RCT) to test for a causal impact in a rigorous way. New initiatives that have been validated by an initial Field Test are prime candidates for a full Experiment.

Adaptive Tactics

The Adaptive Tactic component uses advanced statistics to determine the most effective tactic in response to individual crime events. Feature administrators (analysts or command staff as a department deems appropriate) can input tactics into the system and the Adaptive Tactics component will match these tactics against outcomes to determine which tactic is likely the best for a given scenario. For example, if thefts from motor vehicle are a problem in your department, you can use Adaptive Tactics to determine the best course of action to reduce future thefts. You might enter three tactics: directed patrol for three days after each event, placing flyers on cars warning of the thefts, or even no action. The system will automatically test these tactics and begin recommending the best tactic to you.

Technical Requirements

Hardware Requirements

The HunchLab application is hosted as a multi-tenant Software-as-a-Service (SaaS) application within the AWS infrastructure. Azavea will manage the hosting infrastructure, security updates, and 2nd tier support (Azavea assumes that police departments will prefer to manage direct end-user support). This cloud-based approach enables HunchLab to leverage significant amounts of computing power in an elastic manner, a critical requirement for providing the advanced statistical algorithms the system employs. Replicating a similar environment on-premise would entail a substantial outlay of capital to provide servers that are utilized only in bursts. In order to provide a secure application, we have consulted the FBI's CJIS guidelines to apply as many guidelines as possible in designing our system architecture, including risk mitigation techniques such as encrypting data connections and optional 2-factor authentication.

The application requires network connectivity from the user to the HunchLab service. Bandwidth requirements are modest, as most application assets are cached locally in the browser.

HunchLab 2.0 is hosted within the Amazon Web Services (AWS) infrastructure. AWS provides best-of-breed security and flexibility for building robust and secure SaaS applications.

Software-Related Information (Including Support and Upgrades)

HunchLab's subscription design includes application hosting, updates (fixes and new functionality within the place-based module of HunchLab 2.0), 2nd tier support, and ongoing training resources. This pricing model allows unlimited users and devices to access the application. All support services are coordinated and provided from Azavea's Philadelphia office and will include incident-based and troubleshooting support services by experienced Azavea staff through e-mail or phone during business hours, Monday to Friday, 9am – 5pm, EST (exclusive of designated US federal holidays). Additional support options outside of business hours are available for discussion as needed. Azavea will provide the same level of support during a pilot as would be provided to a long-term client.

Azavea develops HunchLab through an agile Scrum methodology whereby work is planned in 2-week increments. This structure enables us to quickly develop iterative improvements to the application. New functionality and any necessary operating system updates or patches are deployed on a schedule designed to minimize downtime. For instance, most software updates result in about 0 – 15 minutes of downtime. System updates require no work from the client, as Azavea staff manages the deployment process. The application is hosted as a multi-tenant application, so an update by Azavea for the US hosting environment will update all clients hosted within that environment simultaneously.

In order to comply with security requirements, clients are expected to continue to maintain modern, up-to-date web browsers on the devices that will be accessing the system.

Database Requirements

HunchLab does not require a client to have any particular database available. HunchLab does require event data (crimes or calls for service records) to already be geocoded and include basic attribute data such as the date and time of the event (or time range), event classification, unique identifier, etc. More information about the data interface requirements is in the section below.

Data Interface Requirements

HunchLab will consume canonical data sources, such as CAD and RMS systems. The manner in which this data is transferred to HunchLab varies from client to client. A typical process will consist of transferring records to HunchLab as an extension of existing crime mapping and analysis ETL processes. Alternatively, Azavea can configure an upload process that draws data from an ODBC connection to a read-only database view to fetch data that has changed since the last import was conducted. Most agencies schedule this import process on a daily or hourly basis, but HunchLab can also be configured to import changes on a more frequent basis, such as every few minutes. Ideally, historic data is provided for a 5-year period to allow robust predictive modeling.

Event data (crimes, calls for service, etc.) are transferred to HunchLab in a simple CSV format via a secure RESTful API endpoint. Clients can directly push data into this API endpoint or Azavea can support its use. CSV uploads contain column headers and basic attribute values, such as the location and time of an event. Formatted CSV files can also be uploaded directly via the HunchLab administrative UI.

Alternately, Azavea can fetch data from other RESTful APIs, such as those provided by ArcGIS Server. If the endpoint is available via the Internet (with proper authentication), then no on-premise utility needs to be configured as HunchLab can directly fetch updates. If the endpoint is behind a firewall, then the extraction process would be set up within the client environment to push updates to the HunchLab server.

Desktop & Mobile Requirements

Agency staff access HunchLab through a web-browser. As an advanced web-application, HunchLab supports the following major desktop web-browsers and contemporary operating systems:

- Chrome (last two versions)
  - Windows XP or newer
  - Mac OS
  - Linux
- Firefox (last two rapid release versions and supported extended release versions)
  - Windows XP or newer
  - Mac OS
  - Linux

- Internet Explorer (last two versions; IE 10 and 11 as of February 2015)
  - Windows 7 or newer (*Window Vista and older Windows operating systems do not support secure versions of the Transport Layer Security (TLS) protocol 1.2+ within Internet Explorer. To support these older versions of Windows, we recommend the use of Chrome or Firefox (both free for installation) on these machines since they do support these newer versions of TLS.)*

We also support the following major mobile browsers:

- Safari
  - iOS 7 or newer
- Chrome (current version)
  - Android
- Internet Explorer
  - Windows Phone 8.1 or newer

In all cases, the default HunchLab configuration requires operating systems and browsers that support Transport Layer Security (TLS) version 1.2. This requirement is to prevent known attacks against SSL traffic that impact TLS v1.0 and older protocols. Our supported browsers provide the correct version of TLS either automatically or with minor configuration changes (such as checking a box within the settings panel). If an agency is unable to support TLS 1.2 connections, it may necessitate the creation of a separate access point to the HunchLab system, which can be discussed on a case-by-case basis.

User Management

HunchLab can either operate in a stand-alone authentication mode or integrate with existing directory services. In either mode, users are given an application role that provides only the needed functionality to the user. For instance, officers are provided with a viewer role to consume missions. Analysts and IT administrators can be giving access to administrative settings with the application.

To delegate user management to external systems, HunchLab supports the SAML single sign-on (SSO) standard. For instance, a police department's Active Directory system can be published through Active Directory Federation Services (ADFS) to provide a SAML compliant authentication endpoint. The department then simply creates user groups within Active Directory that represent the distinct application roles within HunchLab. Users login to the ADFS service and select "HunchLab" to be redirected into the application with the proper permissions. In this manner, HunchLab maintains no user credentials improving the overall security of the department. As the department deploys advanced authentication techniques, HunchLab benefits automatically by delegating to a centralized authentication service.

Appendix A: Data Guide

Introduction

HunchLab is a predictive policing solution that helps police departments to use their resources more effectively by leveraging advanced forecasts of crime. HunchLab's forecasting methodology fuses many crime theories and data sets into one picture of risk. The system automatically determines how to incorporate concepts such as recent crime events, temporal cycles such as day of week and season, the weather, and geographic locations such as bars and schools to produce a single forecast. The system uses these crime patterns when appropriate without requiring a police department to have a statistician on staff. This approach not only generates robust forecasts of crime but also provides insights into the dynamics of crime patterns.

The forecasting engine uses ensemble machine learning approaches that can incorporate the following crime patterns into a single prediction of criminal risk:

- Baseline crime levels
  - Similar to traditional hotspot maps
- Near repeat patterns
  - Event recency (contagion)
- Risk Terrain Modeling
  - Proximity and density of geographic features (points, lines, and polygons)
- Routine activity theory
  - Offender: proximity and concentration of known offenders
  - Guardianship: police presence (historic AVL / GPS data)
  - Targets: measures of exposure such as population, parcels, or automobiles
- Collective Efficacy
  - Socioeconomic indicators, neighborhood heterogeneity, etc.
- Temporal cycles
  - Seasonality, time of month, day of week, time of day, etc.
- Recurring temporal events
  - Holidays, sporting events, etc.
- Weather
  - Temperature, precipitation, etc.

Types of Information

Event Data

To forecast a space-time event such as a crime, HunchLab requires several years of historic data for the event to build both the outcome variable to be forecasted and several input covariates. At a minimum, HunchLab requires 5 years of crime (event) data for any event being modeled. This quantity of data is necessary to (1) "warm-up" variables that reach into the past up to 1 year, (2) include 3 years of examples to properly model seasonal patterns, and (3) hold back recent data to test accuracy.

This is the only required data set. Reasonably accurate models of crime can be generated with simply this data, but such models do not reveal insights into crime dynamics beyond crime events leading to more crime events.

Event data should be provided for at least the entire area for which forecasts will be used. If data can be provided for a buffer around this region, this can also be included. A buffer of up to 1000m can be useful within the modeling process. A reason to include additional data from nearby areas is that it may increase the overall data volume increasing predictive power. For instance, a small jurisdiction may not incur many violent crimes, but by including violent crimes from nearby jurisdictions more information is presented to the modeling process. Keep in mind that other data sets used in modeling must also be available for the buffered area.

Geographic Data

Geographic layers provide environmental context to the locations at which crimes occur. These datasets change slowly over long periods of time. While HunchLab's analysis is based on a raster format, geographic layers can be provided as points, lines, or polygon layers. HunchLab then builds variables based upon the distance to and concentration of these features. A given geographic layer may be split into multiple layers for the purposes of building covariates. For instance, given a street network for a city, each street segment may be of a different type – highways, highway onramps, residential streets, footpaths, etc. The distance to any street network feature may be a useful feature overall, but building variables for the distance to the nearest highway onramp or footpath may be useful as distinct variables. HunchLab can automatically split geographic layers on 'type' attributes to support this concept.

Implementation staff can easily take static extracts of geographic layers in ShapeFile format for inclusion in HunchLab. This approach requires no integration and therefore incurs no integration fees. Alternatively, a client may desire HunchLab to directly ingest GIS layers from a source such as an ArcGIS Server instance or other web API in GeoJSON format. In such cases, each system from which HunchLab pulls is considered one data connection and the relevant integration fee applies. Ingesting multiple GIS layers from a system does not incur additional fees.

While most geographic data provided by clients is in vector format, HunchLab can also leverage raster layers as variables. For instance, a city may have land cover data in raster format. Such data sets are transformed into a set of covariates and are resampled at the resolution of the HunchLab analysis.

Temporal Data

Temporal data sets provide information about the state of the entire jurisdiction and are considered "global" across the jurisdiction. For instance, a temporal data set may represent when the public school system is in session or the current air temperature. These data sets are provided in CSV format with the relevant time period, variable name, and a numeric value. School being in session may be represented as binary values with a value of 1 when in session and 0 when not in session. The air temperature may be represented as a numeric value in degrees Fahrenheit. Alternatively, the severity of activity between

two feuding gangs may be represented as integers: 0 for no activity, 1 for mild activity, 2 for severe activity.

It is important to realize that any temporal data used in forecasts must be available both historically for several years and for at least 48 hours into the future. The need for future values of the variable necessitates the use of variables that can be manually uploaded far in advance (such as the school schedule) or automation of updates (such as for weather). Temporal data sets that are uploaded into HunchLab manually do not incur integration fees. If instead, HunchLab was configured to automatically pull temporal data from a custom source, then a data connection fee would apply.

Other Variables

HunchLab also leverages variables that are not based upon specific data sets but are, instead, calculated. For instance, the day of the week and day of the month are simply calculated from the date. The moon phase, sunrise and sunset time, and season are other examples of variables calculated in a similar manner.

HunchLab Provided Data

HunchLab has processes in place to automatically manage the inclusion of common data sources if desired by the client. It should be noted that the use of these data sets is not required. For instance, a client may not desire any socioeconomic variables to be used in the forecasts even if academic research suggests it is useful.

Natural Terrain

Elevation data can be automatically loaded into HunchLab. This data set is transformed into several variables that describe the nature of the physical terrain such as the slope and aspect. This data is useful in identifying natural geographic structures that impact settlement patterns.

US Census

The US Census Bureau's American Community Survey provides up-to-date information about the US population based upon a sampled survey of residents. The data is available at the Census blockgroup level. HunchLab can automate the transformation of this data into relevant variables. For instance, this data set can provide measures of the collective efficacy and social cohesion of a neighborhood based upon socioeconomic indicators such as income and the prevalence of renters. The data set also includes information about potential targets of crimes such population density, automobile ownership, and home values.

Weather

Weather data provides a rich source of information about the conditions in a jurisdiction. For instance, seasonal patterns are often found in violent crimes, but these patterns may be more due to the conditions outside (warm temperatures) than the time of year itself. HunchLab can maintain historic weather data and upcoming forecasts automatically for inclusion in models. Variables include such items as the air temperature, humidity, perceived temperature, and precipitation.

Open Street Map

Open Street Map (OSM) is an online, collaborative project to create an editable map of the world. The OSM database includes detailed information about street networks and major points of interest such as schools, libraries, and transportation hubs. HunchLab can use this data if such layers are not readily available from the client.

Thinking About Data Sets

In addition to the above data sets, clients can provide geographic and temporal data for inclusion in HunchLab's models. While more information is often better in building predictive models, a few well-chosen data sets can go a long way to building an accurate and insightful predictive model of crime. We encourage clients to think about this process in an iterative manner as additional data sets can be added over time.

When evaluating a potential data set for inclusion in HunchLab, there are a few key questions to ask:

- Is the data already available? If not, what will be the cost to generate and maintain the data?
  - For instance, a geographic layer that changes infrequently may cost little to maintain while one that changes more often may be burdensome.
- How strongly connected to crime is the data?
  - For instance, if the crime within a jurisdiction drastically changes based upon changes in the student population at a local university, then data sets related to the university are likely quite important.
- Are there synergies between this and other data sets? In other words, does 1 + 1 = 3?
  - The locations of public schools may be useful by itself. The school schedule may also be useful by itself. By providing both school locations and the school schedule, the system can fully identify when and where school may be having an impact. Such related data sets may warrant evaluation as a group.
- Does one set of data represent many ideas?
  - For instance, a city's parcel database may include zoning or land use information that provides information about residential developments, hotels, fast food locations and more.

Ideas for Data Sets

Here are some ideas of data sets that may be useful to include within HunchLab:

- Where people congregate
  - Restaurants, fast food, bars, liquor licenses, nightclubs, places of worship, tourist attractions, movie theaters, exotic clubs
- Where people live
  - University dorms, fraternities, public housing, apartment complexes
- How people move around
  - Bus stops, bus stations, train stations, recreational paths, highway onramps
- Venues for particular types of crimes
  - Pawn shops, retail stores, malls, convenience stores, motels/hotels, ATMs, banks, parking lots, bike parking
- Government buildings
  - Police and fire stations, libraries, post offices
- Problem places
  - Abandoned buildings, vacant lots, foreclosed houses

Additional ideas may be gleaned from the literature reviews of relevant factors for each crime type available for download from the Rutgers University website at http://rutgerscps.weebly.com/publications.html

Appendix B: Application Architecture and Security

Introduction

HunchLab is a web-based server application provided under a software-as-a-service (SaaS) model. While the SaaS model of software deployment abstracts architectural decisions behind a simple client-facing web application, we realize that transparency is necessary within the law enforcement community. This document outlines the architecture of the HunchLab application with an emphasis on application security and availability.

Infrastructure

HunchLab 2.0 is hosted within the Amazon Web Services (AWS) infrastructure. AWS provides best-of-breed security and flexibility for building robust and secure SaaS applications.

Security

AWS data centers maintain strict physical access controls including 24x7, trained security. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. AWS staff members pass criminal background checks prior to employment.

Further, the AWS platform regularly passes third-party evaluations. AWS has achieved ISO 27001 certification and has been validated as a Level 1 service provider under the Payment Card Industry (PCI) Data Security Standard (DSS). AWS annually publishes SOC 1, 2 & 3 audits. AWS is also a FedRAMP Compliant Cloud Service Provider (CSP) with validation at the Moderate level. This validation covers both the regular US regions and the GovCloud region. AWS has been successfully evaluated at the FISMA Moderate level for US federal government systems as well as DIACAP Level 2 for US DoD systems.

AWS Compliance information: http://aws.amazon.com/compliance/

Availability

The AWS platform provides robust services to maintain application availability even in the face of infrastructure failure. Within each AWS region, multiple availability zones allow an application to remain available even with the complete failure of an individual data center. Power and network connectivity systems are designed for redundancy with onsite backup power generation.

The HunchLab application is designed to use multiple availability zones within a region to provide availability even in the face of the loss of a complete zone. For instance, if a client's HunchLab application is hosted within the US East region, client data is replicated between multiple availability zones within the region. Availability zones are independent data centers within the region. The application is designed to survive the complete failure of an availability zone (a complete data center) without manual intervention by Azavea.

Data Residency

Distinct AWS geographic regions allow applications to be deployed to different parts of the world. This allows HunchLab clients to select a region based upon applicable privacy laws. Data placed within a region is not automatically replicated to other regions by AWS.

Clients can select from residency in:

- North America
  - US East (Northern Virginia)
  - GovCloud (US)
- Europe / Middle East / Africa
  - EU (Ireland)
  - EU (Frankfurt)
- Asia Pacific
  - Asia Pacific (Tokyo)
  - Asia Pacific (Sydney)

Additional fees may apply for all data centers except US East (Northern Virginia).

Access

Logical access to the HunchLab AWS hosting account is limited to Azavea personnel working on the application. Access to the infrastructure is granted via 2-factor authentication using individual credentials for each employee. System development and testing occurs in a separate hosting account so that contact with client data is minimized. Client data is not copied outside of the AWS infrastructure without the explicit consent of the client. Statistical models and other diagnostic data that does not include disaggregated criminal justice information (CJI) may be accessed and examined outside of AWS by Azavea personnel for troubleshooting and support purposes.

More details of the AWS platform can be found in the current version of the *Amazon Web Services: Overview of Security Processes* document available for download at http://aws.amazon.com/security/.

Application Security

Azavea has a long history of handling sensitive law enforcement data sets. The new version of HunchLab is delivered as a secure cloud-based subscription service. As we designed this new version, we focused on incorporating security best practices into our development process. While most deployments of HunchLab contain local department data sets that do not technically require compliance with the FBI's Criminal Justice Information Systems (CJIS) guidelines, we are using the CJIS requirements and recommendations to guide our decision-making process and system architecture. Here are some of the security features and policies available within the new HunchLab.

Data Use & Security Agreement

By default, Azavea agrees to solely use the law enforcement data to provide the agreed upon HunchLab service to the department including using the data for system testing, troubleshooting, and lives operations. Separately, Azavea may seek permission to use the data for research purposes that further the product and crime analysis in general. At no time will Azavea hold any claims to the data nor will Azavea use the data for other commercial purposes. Upon written request, Azavea will purge a customer's law enforcement data from its operational systems. Deletion from operational systems will occur within 7 days. The application maintains automated backup files for the last several weeks. Client data will expire out of these automated backups within 28 days from the request. If requested, Azavea will certify that client data has occurred.

Azavea will gladly sign a CJIS Security Addendum as specified in CJIS v5.3 section 5.1.1.5.

Security Awareness Training

Azavea hires technical staff with an eye toward building reliable and secure web applications. Part of the Azavea onboarding process is acknowledgement of company security practices as well as signing a separate agreement regarding confidentiality of client data. Additionally, staff members with access to the HunchLab system undergo biennial training on best practices when dealing with criminal justice information as outlined in CJIS v5.3 section 5.2.

Reliability & Security Incident Management

The HunchLab service is designed to be resilient to failure with redundancy built into the system architecture. Additionally, Azavea has implemented automatic monitoring of system uptime and incident alerts to provide timely resolution of system issues. In the event of a suspected or confirmed security breach, Azavea will proactively notify the law enforcement agency of the breach in a timely manner as specified in CJIS v5.3 section 5.3.2.

System Auditing

The HunchLab system keeps a running system log of activity by users including log-on attempts and information retrieval. These records are retained for at least 365 days. The auditing system is designed to comply with CJIS v5.3 section 5.4. Additionally, Azavea employs AWS services that log the logical access and control of the AWS environment.

Role-based Security

Access to system functionality is restricted based upon security roles. For instance, only a few users need administrative access to the system. This approach reflects the guidelines in CJIS v5.3 section 5.5.2.

Authentication Credentials

HunchLab can delegate credential management to 3rd party directory services such as Active Directory through the SAML standard. In that case, HunchLab assumes that the 3rd party directory service provides a CJIS compliant security model. Additionally, HunchLab can provide a stand-alone authentication system that complies with both the standard authentication and advanced authentication specifications in CJIS v5.3 sections 5.6.2.1 and 5.6.2.2. Our advanced authentication option provides 2-factor authentication using time-based tokens generated locally by mobile applications for mobile devices. Additional costs may apply if Azavea is managing 2-factor authentication on behalf of the client.

Password Management & Login Failures

If operating in stand-alone authentication mode, HunchLab stores user passwords in a salted cryptographic hash format which increases the computing power necessary to reverse engineer a user's password even if our database is comprised. Additionally, to prevent external attacks on user credentials, the system keeps track of unsuccessful login attempts and locks the account for progressively longer periods of time. This policy is recommended in CJIS v5.3 section 5.5.3.

Session Lock

When a user logs into HunchLab, a temporary security token is kept within their local browser memory. HunchLab assumes that devices logging into the system will employ sessions locks or screensavers that meet the guidelines in CJIS v5.3 section 5.5.5.

Data Protection

The HunchLab service is hosted within Amazon Web Services (AWS) data centers. These data centers implement state-of-the-art security practices that protect the physical access to data within HunchLab as recommended in CJIS v5.3 section 5.9. Additionally, AWS continuously monitors their infrastructure against denial of service attacks and penetration vulnerabilities.

Within the HunchLab architecture, Azavea has utilized several security features of the AWS platform to harden the system. For instance, all inbound traffic to HunchLab is encrypted via SSL and terminates at a set of load balancers. These load balancers only allow secure HTTPS traffic with specific versions of the TLS protocol (TLS 1.2+) and specific encryption algorithms (AES) and proxy all traffic to the application. Each component of the application is isolated from all others with only the minimum required network traffic for each server instance granted. This security is enforced as inbound and outbound firewall rules on each server as well as redundantly at the network level.

While the physically secure AWS infrastructure constitutes a physically secure location and therefore encryption is not required, Azavea has decided to encrypt data in transit and at rest as much as feasible. All data in transit within the application is encrypted. Data stored on the elastic block storage devices attached to HunchLab servers and within the AWS S3 service is encrypted at rest. Additionally, data stored in the relational database provided by Amazon RDS is encrypted at rest.

These design approaches seek to conform to CJIS v5.3 section 5.10.

[Note: As of April 2015, there are two pending security features referenced above. We have not yet enabled encryption of temporary files stored on EBS volumes. These files exist for only as long as necessary to process the data and are then deleted. When these transient volumes are released by the application, any remaining data is proactively destroyed by AWS. Second, we have added a caching layer within the application. This presently transmits data within the secure environment without encryption. We are working to encrypt the data being stored within the cache.]

Personnel

Upon request, Azavea will cooperate with the screening of Azavea personnel with access to the HunchLab system in line with CJIS v5.3 section 5.12.

CJIS Policy v5.3 Review

The following review of CJIS Security Policy version 5.2 outlines how HunchLab aligns with these guidelines.

Section	Requirement	Alignment
4.1	Defines Criminal Justice Information (CJI)	The required data set within HunchLab consists solely of crime event data. This data set does not include personally identifiable information. The most sensitive component of the data set is the location of incidence, but this section of the CJIS guidelines exempts property data when it is not accompanied by PII. As such, CJIS does not technically apply.
5.1.1.5	Private contractors are subject to the CJIS Security Addendum when handling CJI.	Azavea will gladly execute agreements in regards to the handling of CJI.
5.2	Security awareness training shall be required within six months of assignment and biennially thereafter for all personnel with access to CJI.	Azavea already conducts new employee briefs on guidelines and responsibilities in handling client data. Specifically to the team responsible for HunchLab, we are implementing focused training to comply with the minimum topics outlined in the CJIS guidelines.
5.2.2	Records of security training	Azavea shall keep records of security training for staff involved in HunchLab projects.
5.3.1	Security events shall be promptly reported.	Azavea shall promptly report security related events to the relevant clients.
5.4.1	Information systems shall generate audit records for specified events.	The HunchLab API logs user interactions that include the event types specified within the CJIS guidelines. Additionally, the AWS environment

		generates audit logs of management interactions with the hosting environment through the use of the AWS CloudTrail service.
5.4.3	Audit monitoring shall be conducted at minimum once a week by designated personnel.	The HunchLab environment generates system alerts upon suspicious activity with a view toward maintaining continuous monitoring of suspicious activity. For instance, increased levels of API requests that fail authentication generate alerts to the HunchLab team.
5.4.5	Protection of audit information from modification, deletion, and unauthorized access.	AWS level audit logs are kept in a secure S3 bucket with modification and deletion access limited to a subset of the HunchLab team. HunchLab API audit logs are kept securely within the hosting environment and end users are prevented from modifying or deleting these records.
5.4.6	Audit records shall be retains for at least one year.	Azavea will retain audit logs for at least one year.
5.5.1	Account management shall be in place to validate system accounts and permissions.	HunchLab client agencies manage user access to the system. Administrative access to the hosting environment by Azavea staff is reviewed regularly with only members of the team granted access.
5.5.2	Access enforcement shall be enforced to limit access to privileged functions.	HunchLab application functions are accessible via role-based system that limits access to administrative features within an organization's account. Additionally, components of the HunchLab application are only granted permissions within the AWS environment for systems that they need access to.
5.5.3	Unsuccessful login attempts shall be limited to no more than 5 consecutive invalid attempt per user followed by an automatic lock on the account for 10 minutes.	This login restriction is in place within the HunchLab application.
5.5.5	Session locks shall be in place to prevent access to the system after	HunchLab assumes that client managed devices will implement screen locks or appropriate

	inactivity.	measures to meet this requirement.
5.5.6	Remote access shall be monitored and controlled.	By its nature a cloud service provides access over an untrusted network. Access to the application is controlled through login requirements. Access to the hosting environment itself is severely limited and requires multi factor authentication and cryptographic keys.
5.6.1	Identification policies should uniquely identify each user or administrator of the system.	All HunchLab users login with a unique identifier. The AWS environment is also managed through unique credentials assigned to each Azavea team member.
5.6.2.1.1	Passwords shall comply with stated attributes.	The AWS environment is managed through unique credentials assigned to each Azavea team member. These credentials include a password (that meets the stated requirement). HunchLab users can have password restrictions assigned to their accounts. Alternatively, if HunchLab is delegating authentication to another system, then that system would enforce such requirements.
5.6.2.2	Advanced authentication is required for publicly accessible services where the authenticity or security of the requesting device cannot be established	HunchLab can either provide 2-factor authentication to end-users directly or can delegate authentication to a client agency to provide a compliant authentication methodology. The AWS environment requires both a password and token to be entered for Azavea staff to access the hosting system.
5.8.1	Electronic and physical media shall be stored within physically secure locations. If not, then the data shall be encrypted.	The AWS hosting environment is a physically secure environment therefore data encryption is not required.
5.8.3	Electronic media shall be sanitized prior to reuse or disposal.	Media within the AWS hosting environment is sanitized before allocation to new customers. Additionally, AWS destroys all media that leaves its data centers for disposal.
5.9	Physically secure locations shall meet stated guidelines.	AWS provides details of its security policies for its data centers. Even Azavea as customers of the service are not permitted physical access to the environment.
5.10.1	The network infrastructure shall	The HunchLab application is comprised of distinct

	control the flow of information between connected systems.	functional units. Each unit can only speak to the other units of the application that are necessary for it to complete its functions. Each server has a firewall that only allows inbound and outbound communication as needed. Additionally, the network enforces traffic controls to specific allowed ports. External access to the environment is limited to a single bastion server accessible only by Azavea.
5.10.1.2	Encryption shall protect data outside of the boundary of a physically secure location when being transmitted or encrypted. Cryptographic modules shall be certified to meet FIPS 140-2 standards	External access to HunchLab is over HTTPS. The application only permits TLS 1.2 due to flaws in earlier TLS versions. The server is configured to use either 128/256bit GCM or 256bit CBC AES encryption and prefers ephemeral key exchange that provides forward secrecy (ECDHE). The application uses Amazon's Elastic Load Balancers to terminate inbound SSL connections. These load balancers do not use certified cryptographic modules unless the application is hosted within the GovCloud environment.
5.10.1.3	Intrusion detection shall be implemented.	AWS manages intrusion detection and abuse of their environment. Additionally, HunchLab logs inbound requests to monitoring servers that provide Azavea staff with a real time view of activity.
5.10.1.5	The metadata derived from CJI shall not be used by any cloud provider for any purposes.	Azavea will not use CJI for any purposes other than to provide this service. AWS also agrees to not use client data for any purposes.
5.10.3.1	Partitioning shall separate user functionality from information management functionality.	The HunchLab application is broken up into discrete segments that separate functionality. For instance, an inbound request for data first arrives at a load balancer, which terminates the inbound SSL connection, parses the request, and then wraps the request in a new SSL connection to pass to the web servers. The web servers then receive the request, validate the user's credentials and query the database for needed data. The database also resides on a separate virtual machine.
5.10.3.2	Virtualization shall be	Firewalls are in place that restrict access to each

	implemented to isolate machines.	machine within the environment. For instance, the load balancers may not directly communicate with the database server. Requests from the load balancers to the web servers and from the web servers to the database server are encrypted. Requests from all servers to the S3 object store are all enforced as encrypted. Log files on each machine are centrally aggregated for monitoring.
5.10.4.1	Patches shall be maintained.	Azavea maintains a staging environment to validate updates to software. The application utilizes OS releases that are currently supported with security patches. OS security patches are applied upon each deployment of the software via golden images for machines. Additionally, Azavea updates other software packages on a regular basis based upon the severity of the patch.
5.10.4.2	Malicious Code Protection	HunchLab utilizes only Linux based software. It is atypical to run antivirus software on such systems due to the security design of the systems. Additionally, the hosting environment is designed for the rapid replacement of server instances based upon golden images. For instance, no persistent data is stored on web application servers. Upon every deploy of an update, the existing servers are destroyed and replaced with new servers running from a clean image. This approach eliminates the likelihood of an infection of maintaining itself.
5.12.1	Personnel will have fingerprint-based record checks.	Azavea is happy to have relevant staff cleared through these processes.
5.12.2	Upon termination, access to CJI shall be terminated immediately	Azavea maintains a checklist of termination practices, which includes removal of access to the HunchLab environment.

Application Architecture

Overview

The HunchLab application is hosted as a multi-tenant SaaS application within the AWS infrastructure. The application leverages a broad array of open source projects including operating systems, application frameworks, and statistical packages. Further, the application leverages AWS-specific technologies to provide scalability, redundancy, and security. Finally, the application is architected into discrete tiers allowing the logical separation of components.

Open Source Technologies

The application consists of a client-side, standards-based web GUI application implemented in JavaScript using the Angular JS framework. This GUI application speaks to a set of RESTful APIs implemented in the Django web application framework with data persistence provided by an AWS-managed PostgreSQL database with geographic queries supported by the PostGIS extension. Additionally, the system uses Azavea's open source GeoTrellis framework for high performance geographic processing and the R framework for state-of-the-art machine learning algorithms.

Amazon Web Services Technologies

The HunchLab SaaS application was designed to take advantage of the breadth of AWS services to provide a secure and scalable application. The application uses the following AWS technologies:

- Route 53
  - DNS for the hunchlab.com domain is managed through the distributed and redundant Route 53 service.
- Virtual Private Cloud (VPC)
  - VPC allows the isolation of application components on individual subnets, enforces network-level traffic rules, and provides both inbound and outbound firewalls.
- Elastic Load Balancing (ELB)
  - The SSL encrypted web traffic for the application is terminated by elastic load balancers which provide secure management of the signed HunchLab SSL certificate and performance under increased application loads.
- Elastic Compute Cloud (EC2)
  - Servers provided by EC2 are used for the web application, database, and machine learning tiers of the application. Many AWS services utilize EC2.
- Elastic Block Storage (EBS)
  - EBS volumes back the root partitions of EC2 instances and are used to store client-specific data.
- Relational Data Store (RDS)
  - RDS provides a managed relationship PostgreSQL database to HunchLab. The RDS instance is configured for real-time replication and automatic failover between availability zones.
- Simple Storage Service (S3)

  - Additional application artifacts are stored in the S3 service using the server-side encryption option. These artifacts include data sets undergoing processing, analytic models and results, and backup files.
- Glacier
  - Long-term backup archives are hosted in the Glacier service.
- ElastiCache
  - An in-memory application cache is provided by the Redis functionality of ElastiCache.
- Simple Workflow Service (SWF)
  - Machine learning and prediction processes are managed via the SWF service allowing distribution of tasks among a cluster of compute instances that scales to meet client needs.
- CloudWatch
  - CloudWatch metrics and alarms are used to scale application resources to meet demand and to notify Azavea staff of failures.
- CloudFormation
  - CloudFormation is used to securely manage and update the application stack with discrete application components isolated from one another and designed to automatically scale to meet user load.
- Identity and Access Management (IAM)
  - IAM is used to provide individual credentials to Azavea staff tasked with supporting the application. IAM security policies require the use of 2 factor authentication tokens when interacting with the AWS infrastructure. Additionally, IAM security roles are used within the application stack to provide credentials to application components.
- CloudTrail
  - CloudTrail provides audit logs of interactions with AWS management commands. These logs are stored securely within S3.

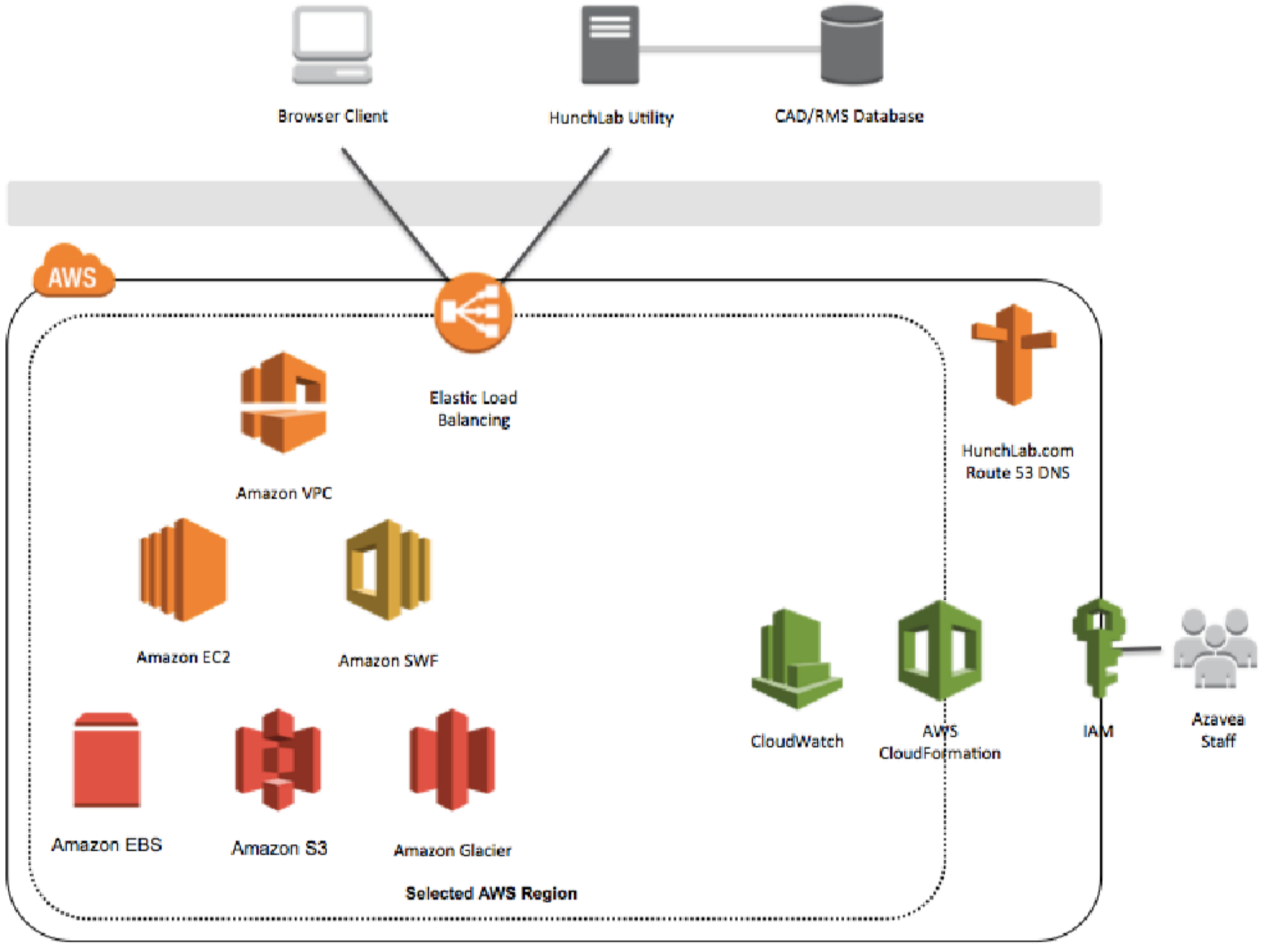

Application Components

The HunchLab SaaS application is designed as a set of loosely coupled components that work together to service the user. The main components of the application include:

- Client-side
  - Browser-based application
    - JavaScript application that provides the graphical user interface
  - HunchLab data upload (varying formats based upon client needs)
    - Data integration utility for crime data
- Server-side
  - Web application tier
    - Serves static files and provides RESTful APIs consumed by the browser application and integration utility
  - Geographic processing tier
    - Conducts geographic processing to support requests from the web tier
  - Machine learning tier
    - Batch processing for creating statistical models and generating crime predictions

- Persistence tier
  - Relationship persistence for the web tier and file persistence for objects shared among tiers

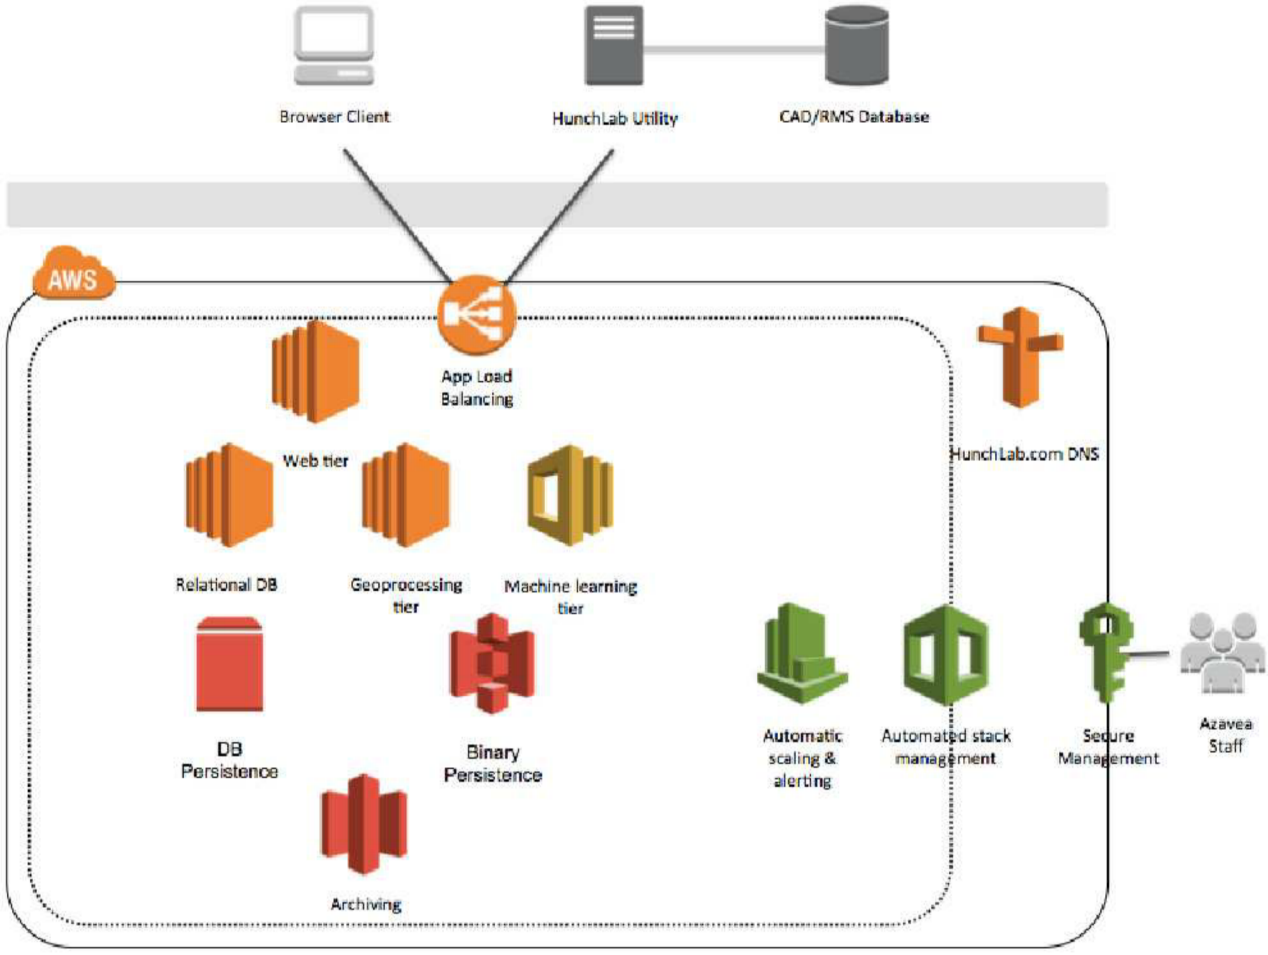

Appendix C: FAQs

What are the possible methods for uploading data to your product?

Event data can be uploaded directly via the graphical interface in a simple CSV format. Alternatively, the data uploads can be automated via submission to our RESTful API. IT staff can either configure this process or Azavea can set it up. Another option is for HunchLab to fetch data from other APIs such as the ArcGIS Server API.

More detailed information is available in the "Interface Requirements" section of our response.

What type of crime does the software predict?

Our current functionality in HunchLab 2.0 focuses on forecasts of specific crimes in specific geographic locations. The number and types of crimes can be configured by an individual police department. For example, an analyst may configure a crime model within HunchLab to forecast residential burglaries. HunchLab would then automatically produce predictions (count expectations) of residential burglaries in each raster cell (configurable, but each cell is often 100-250m in size) for the upcoming hours. The temporal granularity can range from a one-hour block of time to one shift.

Individual police departments would most likely configure several crime models within HunchLab based upon the manner in which they organize themselves. For example, you would likely create a model for each major crime type, as well as additional focus crimes that the police department is looking to address. The system generates separate predictions for each crime model that is configured. These separate predictions are then combined to create target areas based upon the crime weights set by the police department.

Because our system leverages more than just crime data, it is more adaptable to forecast various types of crime. The system automatically learns what information it has access to best forecasts a given type of crime. So far we have tested the system with the following crime types:

- Homicide
- Aggravated Assault
- Simple Assault
- Robbery
- Residential Burglary
- Commercial Burglary
- Theft Person
- Thefts of MV
- Thefts from MV
- Motor Vehicle Accidents
- Bike Thefts
- Vandalism
- Rape

How is the information viewed by police staff?

As a web-based application, HunchLab mission areas can be viewed through contemporary web browsers on devices ranging from desktop computers and laptops to tablets and MDTs. For viewing the data in the field, the simplest dissemination technique is a printable PDF report outlining mission areas for the shift. For departments with mobile broadband and GPS-enabled MDTs, smartphones, or tablets HunchLab also provides a location-based service called Sidekick. Sidekick provides officers with crime predictions about their current location, notifies them of mission objectives, and supports measurement of the dosage of field tactics to address crime problems. Access to specific system functionality is restricted based upon security roles.

What data /sources are used to predict crime?

HunchLab 2.0 can use both police data and other data sets to generate crime forecasts. The only required data set is the crime event data itself (the outcome being forecast). Azavea tests our statistical models with just this data set being available to ensure that we can produce accurate forecasts if this is the only data that a specific police department has available.

By including other data sets in HunchLab's modeling process, a diverse picture of criminal risk is portrayed. Some data sets we manage on behalf of clients. For instance, we automate the use of national holidays and weather within our forecasts. Other data sets might be procured by the police department from another government agency (example: business permits for locations of restaurants or liquor establishments) or developed internally (example: gang member residences). These additional data sets fall into two main categories:

1. Geographic data sets containing points, lines, and polygons
2. Temporal data sets, such as school schedules, social events, or weather data

How is GIS information obtained?

HunchLab provides default options that satisfy the application's core GIS requirements. Additional GIS services can be used within the system to customize the application for use by the client or to enhance the analytic process.

Supported Options By Requirement

Basemaps provide the map upon which HunchLab displays data. HunchLab supports:

- Standard Options
  - ArcGIS Online basemaps
    - Previews: http://www.arcgis.com/home/webmap/viewer.html?useExisting=1
    - Choices
      - Topographic
      - Imagery
  - Stamen Toner basemap
  - Azavea's custom-styled OpenStreetMap basemap
- Custom Options
  - ArcGIS Online basemaps

        - Clients can publish custom managed basemaps to ArcGIS Online for consumption within the HunchLab application.
    - Mapbox tiles
        - Clients can push custom tile sets to their Mapbox account for display within HunchLab.

Contextual Data Sets provide situational data layers for use in analysis; this is an optional requirement. HunchLab supports points, lines, and polygon datasets in these formats:

- Static Sources
    - Shapefile format upload
- Live Sources
    - ArcGIS Online
    - Public ArcGIS Server REST APIs

Geographic areas are loaded into the system for query and analysis purposes. These boundaries specify areas such as the jurisdictional boundary, counties, cities, police divisions and districts, and even census geographies. Geographic boundaries are displayed in a hierarchy to a user. For instance, Division 1 may contain Districts A and B. Ideally, these geographic levels would nest perfectly within one another. Geographic areas are uploaded in ShapeFile format.

Is the information available outside city network?

HunchLab 2.0 is hosted within the Amazon Web Services (AWS) infrastructure and is available (with proper authentication) over the Internet. If this is problematic, Azavea could discuss addressing the needs that you have in this regard.

AWS data centers maintain strict physical access controls including 24x7, trained security. Authorized staff must pass two-factor authentication a minimum of two times to access data center floors. AWS staff members pass criminal background checks prior to employment.

Logical access to the HunchLab AWS hosting account is limited to Azavea personnel working on the application. Access to the infrastructure is granted via 2-factor authentication using individual credentials for each employee.

System development and testing occurs in a separate hosting account, so that contact with client data is minimized. Client data is not copied outside of the AWS infrastructure without the explicit consent of the client. Statistical models and other diagnostic data that does not include disaggregated criminal justice information (CJI) may be accessed and examined outside of AWS by Azavea personnel for troubleshooting and support purposes.

Please provide success story examples from organizations that have similar crime and population demographics.

Since HunchLab 2.0 is new, Azavea does not yet have published success story documents. The theoretical concepts used within HunchLab 2.0 are not new, however, and their success has been studied and documented by many academics. At the end of this section, Azavea has provided links to some research papers that we and others in the field have written on Risk Terrain Modeling, near repeat forecasting, and other methodologies that have been operationalized in HunchLab. Individually, these methodologies have demonstrated their effectiveness in a number of police departments in North America and Western Europe, but research is still ongoing.

For HunchLab 2.0, Azavea views success as two main components. The first component deals with the accuracy of the predictions themselves. This component we are measuring for various crime types in an ongoing fashion and can make claims about. The second component deals with the effectiveness of a predictive policing tool in terms of crime reductions. It's appealing to be able to include a glowing report on the effectiveness of predictive policing tools for deterring crime as part of our realized benefits, but these results are very difficult to prove. Software, after all, does not prevent crime on its own, but it can have a substantial preemptive impact, depending on how it is used by the agency that deploys it. For example, the same tool might be deployed at two different law enforcement agencies in the same geographic region. If the culture of the first agency is very data driven, they will likely adopt the tool to great success. If the command structure at the second agency does not value the tool or use it extensively, it is likely to have limited impact. This fact is a key reason why we encourage that agencies try tools operationally through a pilot to prove that there is a culture fit before committing to a product long-term.

Here are some publications of interest on HunchLab's proven crime modeling techniques:

Near repeat pattern analysis

Haberman, CP & Ratcliffe, JH (2012) The predictive policing challenges of near repeat armed street robberies, Policing: A Journal of Policy and Practice.
http://jratcliffe.net/papers/Haberman_Ratcliffe_2012_Predictive%20policing%20challenges%20of%20armed%20street%20robberies.pdf

Ratcliffe, JH & Rengert, GF (2008) Near repeat patterns in Philadelphia shootings, Security Journal. Volume 21, issue 1-2: 58-76.
http://jratcliffe.net/papers/Ratcliffe_Rengert%20(2008)%20Near%20repeat%20patterns%20in%20Philadelphia%20shootings.pdf

Risk Terrain Modeling

Heffner, J. (2013). Statistics of the RTMDx Utility. In J. Caplan, L. Kennedy, and E. Piza, *Risk Terrain Modeling Diagnostics Utility User Manual (Version 1.0)*. Newark, NJ: Rutgers Center on Public Security.
http://www.rutgerscps.org/software/index.html

Additional resources (publications, software, manuals):
http://www.rutgerscps.org

Seasonality
Wilpen Gorr , Andreas Olligschlaeger , Yvonne Thompson (2003) Short-term Forecasting of Crime, International Journal of Forecasting 19
http://forprin.dev.zoe.co.nz/files/pdf/Gorr_Olligschalger_and_Thompson,_Short-term.pdf

Appendix D: Experience and Qualifications

Brief History of the Firm and Experience with HunchLab

Azavea is an award-winning geospatial software design and development company based in Philadelphia. The firm was organized in 2000 to create technologically advanced solutions for web and mobile geospatial data visualization and analysis. Azavea is a certified B Corporation, a for-profit corporation with a social mission. Our mission is to apply geospatial technology for civic and social impact while advancing the state-of-the-art through research. Azavea provides a range of services that include:

- Web and mobile software development
- User interface and experience design
- Mapping and spatial analysis
- High performance computing
- Spatial data mining and modeling
- Research and development

The firm has designed and implemented geographic data applications for a variety of domains including: crime analysis, public safety, economic development, elections, urban forestry, humanities and land conservation.

Technology and Partners

Azavea's developers work with a broad range of tools and have particularly strong backgrounds with the .Net, Java, Python, Django and Scala frameworks. The firm has also established a number of strategic partnerships to enhance our capabilities.

Azavea is a member of the Amazon Partner Network, which provides us with a higher level of support for the HunchLab infrastructure we host in the Amazon Web Services environment.

Azavea is an Esri Business Partner and has several years of experience with development and deployment on the ArcGIS platform with dozens of applications implemented on the ArcGIS Server, ArcGIS.com and ArcIMS products. Azavea was named ESRI Business Partner-of-the-Year or Foundation Partner-of-the-Year in 2006, 2007 and 2010. In addition, Azavea is a Microsoft business partner with substantial experience developing the .Net Framework, SQL Server and Windows Server platforms.

In addition to commercial toolkits, Azavea staff is experienced creating web software solutions that use online API's such as GoogleMaps, Bing Maps, ArcGIS Online and OpenStreetMap. The firm also works

with a range of open source tools that accelerate and lower the cost of our software development work. In particular, Azavea has a great deal of experience with creating solutions that bring together the strengths of both commercial and open source toolkits to create high quality and visually attractive applications. The firm not only has experience with open source solutions, but also contributes to them, including significant contributions to OpenLayers and PostGIS.

User Experience Design

Azavea takes great pride in the development of user interfaces that are simple, easy-to-use and are crafted for the specific purpose at hand. Our talented developers and designers work with each client to develop applications that aren't simply functional, they are simple and beautiful.

Commitment to Community

Azavea is committed to working on projects with a strong social value component. Each of Azavea's projects, products and pro bono engagements showcases this commitment. We seek out projects that enhance communities, foster economic development and improve decision-making. Further, we perform research to advance the state of the art.

Azavea R&D

Azavea has an active research and development program through which the firm invests substantial resources toward the development of new solutions and techniques. Each employee is encouraged to develop a personal research project that will both engage the employee and extend the capabilities of the organization. Current research projects include: an effort to apply genetic algorithms to generate transit routes; UI/UX design experiments with Google Glass; application of the OpenCV computer vision framework to recognizing trees in Google Street View; better PostGIS import tools; Emacs integration with Django; and an exploration of machine learning algorithms for space-time forecasting. While not all of these research projects results in measurable commercial success, they are an important part of a culture at Azavea that encourages and takes pride in innovative applications of geospatial technology.

The HunchLab Project

Current law enforcement trends reflect an interest in predictive policing, increasing adoption of cloud services, pervasive location information, and the proliferation of mobile devices. These trends are emerging in extremely constrained budget environments for many agencies and communities. HunchLab combines geospatial data with the modern capabilities of machine learning algorithms to help police departments make effective use of their limited resources in the fight against crime.

The HunchLab project began in 2005 as a prototype to detect localized spikes in crime activity. Azavea then secured research and development funding from the National Science Foundation to expand this prototype into a commercial product. Development of this first version of HunchLab proceeded from 2008 through 2011 and included early versions of basic forecasting capabilities. In 2013, we developed a new statistical approach to forecasting crime that could include multiple data sets and crime theories in one unified picture of risk. This new statistical approach was combined with a new design to form the HunchLab 2.0 application and previewed publicly in September 2013.

From its earliest iteration, HunchLab's crime risk forecasting – or "predictive policing" – techniques have been based upon published academic research that has looked at an individual aspect of crime patterns. For example, Azavea software developer and data scientists worked with Professor Jerry Ratcliffe at Temple University to create a daily risk forecast in HunchLab for burglaries, shootings, and other crime based on his near-repeat pattern research. Police officers have understood for many decades that, for some crimes, the risk of being a repeat victim is quite high. In other words, if someone is a victim of a burglary, there is actually a significant chance that they will be a repeat victim in the weeks after the initial crime. But Ratcliffe and his colleagues discovered something even more interesting. Not only is there an elevated risk that someone will be a repeat victim, but that the risk of their neighbors becoming a victim is also higher for a few weeks after the initial crime.

In addition to the near-repeat pattern phenomenon, Hunchlab includes concepts such as the Risk Terrain Modeling research being published by Rutgers University. Risk Terrain Modeling describes geographic location through correlated geographic features such as bars, schools, and transit stops. This permits the forecasting of crime locations not because crimes occurred there yesterday, but because the social and environmental conditions are ripe for crimes to occur there in the near future.

HunchLab includes each crime theory by deriving individual sets of variables that represent the underlying concepts. For example, Risk Terrain Modeling may be represented by measuring the distance to the nearest bar and the density of bars in each raster cell. Near-repeat patterns may be represented by measuring the amount of time since the most recent crime occurred in each raster cell. These sets of variables are then passed into the modeling process, which determines the useful theories for a given crime type. The system also determines how the theories interact. For example, if the near-repeat pattern effect is stronger in areas with lots of historic crimes, the system can use that information to enhance the forecast. If assaults frequently happen on Friday evenings near bars, the system can similarly model that effect. HunchLab incorporates machine learning concepts to help the software "think" like a crime analyst by imitating years of experience drawn from a police department's own data.

List of Related Clients

HunchLab is an ideal solution for municipal and regional law enforcement agencies. Early versions are deployed in the City of Tacoma and Pierce County, Washington, the Philadelphia Police Department, and the Northwest Ohio Regional Information System. The first class of HunchLab 2.0 pilots include the City of Philadelphia, New Castle County, Delaware, Lincoln, Nebraska, and Greensboro, North Carolina. Additional pilots are planned or underway across North America and Western Europe, but Azavea is often under nondisclosure agreements during these projects. Each deployment is agency-wide and may involve many dozens of crime analysts and hundreds of patrol officers.

In addition to our HunchLab product, Azavea has designed and implemented web and mobile geospatial solutions for a wide range of domains. We have deployed these applications for government organizations at the local, regional, and national levels. A few highlights of these projects include:

Lincoln Police Department: HunchLab 2.0

Lincoln, Nebraska

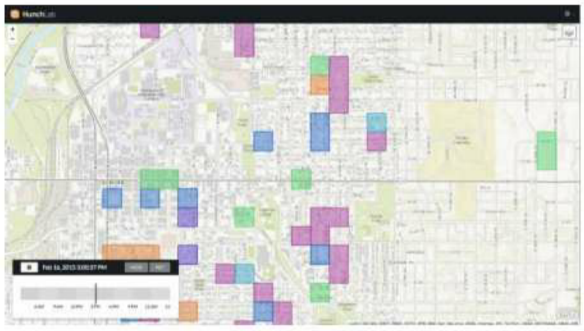

The Lincoln Police Department routinely uses a range of community policing tactics other than responding to individual incidents, such as: targeted saturation patrol, bicycle and foot patrol, undercover/plainclothes/surveillance operations, educational presentations, and coordination of efforts with other government or human service agencies. The Lincoln Police Department has shared crime data with Azavea to develop a HunchLab 2.0 instance that underscores HunchLab's ability to meet community policing initiatives.

Firearms Analysis System (FAS)

Bureau of Alcohol, Tobacco, Firearms and Explosives

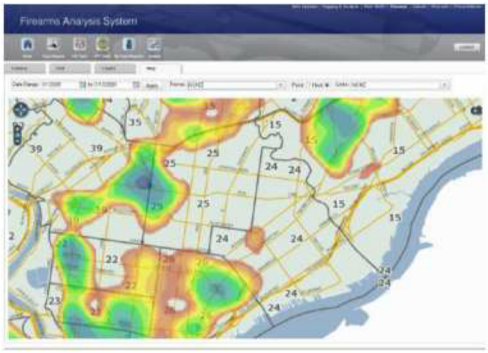

The Firearms Analysis System (FAS) was developed to reduce the amount of time required to trace a firearm seized in a crime and to provide greater analytical ability to look at all such crime guns seized in Philadelphia. The FAS serves both as a workflow / data entry system and as a mapping, charting, search, and reporting system. The Philadelphia Police Department worked closely with the Philadelphia Division of the Bureau of Alcohol, Tobacco, Firearms and Explosives (ATF) and Azavea to develop the application.

Priority Places: http://priorityplaces.ashevillenc.gov/

City of Asheville, NC

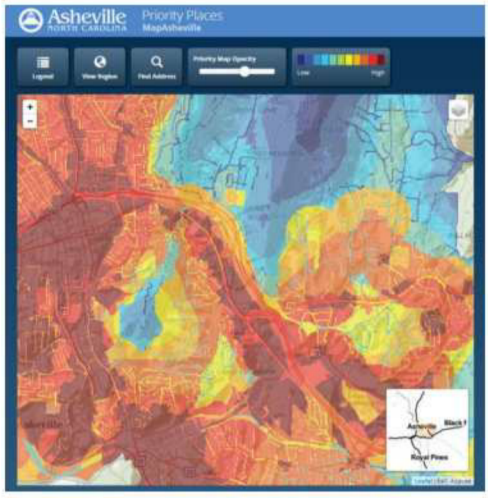

Priority Places was developed to enable citizens and entrepreneurs to pick great sites for facilities, activities and businesses in Asheville, North Carolina. Visitors to the website can select a series of "decision factors" to create a siting scenario, assign weights to each factor and then generate a "heat map" showing the best locations relative to the rest of the City. Users can then select individual property parcels to find detailed information about each lot. They are able to select from multiple base maps and overlays, adjust the color ramps and transparency, and select from different scenario templates to customize the maps to meet their needs.

Wellbeing Toronto: http://map.toronto.ca/wellbeing/

City of Toronto, ON

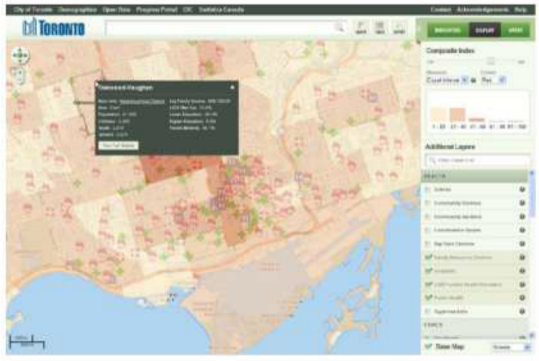

Wellbeing Toronto was developed for the City of Toronto, Ontario, Canada to help users learn more about the City's 140 diverse neighborhoods. The socioeconomic indicator data presented in the application is important to the public when searching for a new home, planning a business, investigating crime trends or learning more about a neighborhood of interest. Users are able to combine and weight the importance of indicators ranging from economics to health to transportation or the environment in order to create customized mapping outputs that meet their individual needs.

PhillyStormwater: www.phillystormwater.org

Philadelphia Water Department

Stormwater billing based on impervious surface area both distributes the cost of managing stormwater more equitably and creates incentives for property owners to implement best management practices for mitigating stormwater runoff on-site. The Philadelphia Stormwater Billing application enables Philadelphia property owners to explore environmentally friendly strategies that will improve the urban environment and enable them to save money on their monthly water bills.

Water Infrastructure Systems Data Manager (WISDM):
US Army Corps of Engineers

Azavea had developed a prototype Water Infrastructure Systems Data Manager (WISDM) tool, a web-based software and services solution for the U.S. Army Corps of Engineers Institute for Water Resources (IWR). The application is based on Azavea's GeoTrellis high performance geoprocessing engine. WISDM will enable real-time interactive spatial analysis of large data sets in support of collaborative sustainable water resource development. IWR has assembled a range of watershed metrics that can contribute to these priority maps.

Appendix E: Unique Features

Produces Forecasts Fusing Disparate Data Sets

To our knowledge, HunchLab is the only software-as-a-service (SaaS) solution on the market that can leverage multiple types of data to arrive at a single picture of criminal risk. Most crime forecasting systems are limited to only leverage crime event data. Some firms can use additional data sets within statistical models but do so within custom engagements and are not providing an off-the-shelf SaaS solution at an affordable cost.

Automatically Adapts to Forecast New Crime Types

The machine learning algorithms used within HunchLab are unique in that they do not rely on assumptions that a simple pattern (near repeats, self-excitation) is the best way to forecast all crime types. The modeling process is presented with a diverse set of patterns that may be relevant to forecast a given type of crime and then automatically learns what patterns matter both individually and as arbitrarily complex interactions. This quality means that HunchLab is often able to accurately forecast new types of crime without manual intervention or additional enhancements.

Provides Crime Pattern Insights

Crime forecasting systems that only use crime data are limited to creating forecasts that explain crime based upon past crimes. By including other data sets and presenting different patterns on an even footing, the system determines what is most useful to produce accurate forecasts. Beyond accurate forecasts, however, HunchLab provides insights into the dynamics of crime represented within these models.

This capability is unique across products on the market.

Delivers Accurate Forecasts with Validated Measurements

When HunchLab builds forecasting models, it automatically tests variations of the models and measures accuracy against a held-out data set on each client's data. Compared against six base-line models representing variations of what an analyst may produce, HunchLab can prioritize risky areas substantially more accurately. The system can also provide the expected level of accuracy to the client so that operational decisions can be made appropriately.

Aligns Police Deployments with the Community Impact of Crime

HunchLab enables police departments to systematically allocate police resources based upon the harm cause by crime in communities. For instance, clients tell HunchLab how much harm a robbery causes in comparison to a burglary. The system then uses this to make intelligent decisions of how to allocate the resources that are available to reduce the overall harm of crime to the community.

No other software product has this capability.

Supports an Open and Transparent Approach

The original research conducted by Azavea to build HunchLab has been documented publicly and is available in 3rd party open source packages such as the R project. As such, HunchLab is unique in that it is the only crime forecasting package on the market that can be recreated with open source technologies without violating patents. This approach both reduces long-term lock-in and enables 3rd parties to examine our methodology.

Provides a Highly Available Subscription Service

HunchLab leverages Amazon Web Services to deliver advanced forecasting as a simple SaaS solution. This provides distinct advantages in comparison to any on premise solution. For instance, the application always operates across two or more distinct data centers separated by several miles. In the event of a natural disaster or other loss of an entire data center, the system will remain available without intervention by Azavea. The application is also entirely self-healing. For instance, if a web server no longer responds to requests, it will automatically be replaced with a fresh server without intervention by Azavea. Finally, the system can leverage the elastic nature of the cloud to use substantial quantities of computing power to meet changes in demand without requiring the investment in physical servers.

Most crime analysis packages are not designed to leverage the distinct nature of the cloud to provide these robust capabilities and, instead, run on a single server prone to hardware and software failures.

Secures Data Robustly

In designing the architecture of HunchLab, security is of forefront concern. HunchLab benefits from the leading best practices of the Amazon Web Services environment. The application supports robust encryption of data in transmission to the application via FIPS 140-2 recommended ciphers. Additionally, the application encrypts requests in transit within the cluster and when data is stored permanently either in flat files or within the database.

Delivers a Cost Effective Solution

Because of the SaaS design of HunchLab, it delivers more advanced analysis at a price point lower than the custom solutions offered by consulting firms or the simple forecasting packages offered by other predictive policing vendors. We have found HunchLab to be about 40% cheaper than quotes from our closest competitors and an order of magnitude cheaper than custom solutions.

Informed by Research and Academic Collaboration

Azavea is a research-driven organization. The first version of HunchLab was developed with the help of a research grant from the National Science Foundation. In addition to former and ongoing research projects with Temple University, Azavea wrote the statistical models powering the Rutgers University Risk Terrain Modeling Diagnostics Utility. The research and development from each of these projects helped inform the development of HunchLab 2.0. We believe in contributing to the crime analysis domain as a whole while building our business.