ABSTRACT

The New York City Police Department's Lower Manhattan Security Coordination Center integrates data from a variety of sources, including sensors (cameras, license plate readers, and environmental detectors) and records (arrests, complaints, summonses, 911 calls, etc.). Analyzing this data to inform decision making has required the development of several coordinated processes. These processes are leading to increased efficiency and effectiveness as well as improved situational awareness for senior leadership.

INTRODUCTION

The INFORMS analytics section defines analytics as the ''scientific process of transforming data into insight for making better decisions.'' Police agencies put analytics into action every day to aid in their deployments and investigations, achieving increased efficiency and effectiveness (for many examples see Langworthy [1999] and also the discussion in chapter 6 of National Research Council [2004]). Recently, the data available to police agencies has expanded dramatically, while also becoming more easily accessible to analysts, mirroring a similar private sector trend (Davenport and Harris, 2007). In this context, police agencies must develop analytic processes to ensure that relevant pieces of information are extracted and prioritized from the large volume of data available.

In this article, we will describe the analytic processes that our organization, the New York City Police Department (NYPD), has developed to handle the large amount of data aggregated at a specialized operations center called the Lower Manhattan Security Coordination Center (LMSCC). It is our hope that by sharing the lessons learned in building these processes we will inform the development of operations centers across both the public and private sectors. In the next section, we describe the context for our analytic processes, including the organizational structure and the history of analytics at the NYPD. Next, we provide an overview of the data sources integrated at the LMSCC, and then discuss the analytic processes we have developed to manage these data sources. We document examples of the system and its processes in action in the subsequent section. We then discuss some of the benefits and challenges associated with applying analytics at the NYPD and at the LMSCC.

ANALYTIC CONTEXT

Organizational Structure

The NYPD is the largest municipal police department in the United States (Reaves, 2010), employing approximately 34,400 officers and 15,800 civilians (New York City, 2012). With an annual budget of approximately $4 billion, the NYPD's jurisdiction encompasses the five boroughs of New York City. The NYPD divides the city into 77 precincts; each precinct includes separate patrol and detective commands. Additionally, separate areas of responsibility have been delineated for transit and public housing; these commands evolved from the merger of independent transit and housing police organizations with the NYPD in 1995. Many specialized units also exist, distinguished by their capability (e.g., aviation, harbor, and mounted units) and responsibility (e.g., intelligence, narcotics, auto crime, and organized crime).

One of these specialized units is the Counterterrorism Bureau (CTB), created in the wake of the terrorist attacks of September 11, 2001. The CTB was the first municipal counterterrorism law enforcement bureau in the United States, and is charged with protecting the city from the unique and singular threats New York City faces from both domestic and international terrorism. The tools and tactics CTB uses to accomplish this mission include investigation, high-visibility deployments, infrastructure protection, public-private partnerships, intelligence analysis, and surveillance, among others.

The CTB is an integral piece of the national counterterrorism effort. The US Department of Homeland Security believes that ''the building, sustainment, and delivery of [core capabilities essential to achieving . . . a secure and resilient Nation] are not exclusive to any single level of government . . . but rather, require the combined effort of the whole community'' (FEMA, 2014). In other words, the CTB's work at the municipal level complements the efforts of the federal government. In fact, federal

Analytics in Action at the New York City Police Department's Counter-terrorism Bureau

E.S. Levine and J.S. Tisch

New York City Police Department
evan.levine@nypd.org,
jessica.tisch@nypd.org

APPLICATION AREAS: Homeland Defense and Civil Support

OR METHOD: Pattern Recognition, Categorical Data Analysis

grant programs provide much of the funding for the CTB, including the Urban Area Security Initiative, Port Security Grant Program, and the Transit Security Grant Program administered by the US Department of Homeland Security's Federal Emergency Management Agency, and the Securing the Cities program administered by the Domestic Nuclear Detection Office.

The LMSCC, a facility located a few blocks from the World Trade Center and staffed 24 hours a day, seven days a week by NYPD officers, is part of the NYPD's CTB. The LMSCC is the central aggregation point for the information gathered by the CTB's citywide sensor network.[a] The officers assigned to the LMSCC use the available data to inform decisions with regard to counterterrorism operations and investigations. In essence, their mission is to perform analytics on the NYPD's data sources.

History of Analytics at the NYPD

In some organizations, establishing an analytic culture is the most difficult part of becoming an analytically mature organization (Davenport et al., 2010). In most police agencies, some of the groundwork for an analytic culture has already been laid because police officers are trained to use deduction, structured thinking, and hypothesis testing to investigate crimes. These skills happen to be key ingredients of an analytic culture and are necessary to take advantage of the large amount of data now flowing into police agencies.

At the NYPD, some additional advantages in establishing an analytic culture are present because of preexisting analytic processes. CompStat, the NYPD's structured methodology for examining crime data, is a well-known example of one of these processes (Henry, 2002). CompStat began at the NYPD; similar analytic processes have been implemented at many police agencies and other government organizations around the world (Weisburd et al., 2003). A recent study demonstrates that CompStat is a likely contributor to the decline in New York City's crime rate since 1990 (Zimring, 2012).

When it was first implemented, the goals of CompStat were to ensure the accountability of precinct commanding officers, improve performance, and share information (Silverman, 1999). CompStat accomplishes these goals via regular meetings during which the commanding officers and staff of each precinct are required to think analytically about crime in their area of responsibility. During the meeting, the weekly crime statistics are presented and discussed, including geographic and temporal trends and patterns.

An additional piece of evidence for a preexisting analytic culture at CTB is the Bureau's development and publication of a methodology to assess by risk potential terrorist targets (NYPD, 2009a). Similar methodologies are common across the homeland security community; see, for example, the Department of Homeland Security's tiered list of critical chemical infrastructure (Beers, 2013). These methodologies are particularly useful when speaking to stakeholders without analytical backgrounds.

LMSCC DATA SOURCES

The LMSCC integrates data from two categories of sources: sensors and records. Sensors are devices through which the NYPD makes a direct observation; they typically send a continuous stream of information whenever the sensor is operational. Records, on the other hand, are documents that are created and maintained by the NYPD or some other organization. After records are created they are sometimes revised or updated. Information from both sensors and records is useful for informing decision making.

Sensors

The NYPD has deployed a networked assortment of sensors of various types, including cameras, license plate readers, and environmental sensors.

Cameras. Closed-circuit television cameras are an increasingly common policing tool used to provide situational awareness and forensic support for both law enforcement and counterterrorism purposes. For example, Chicago's Office of Emergency Management and Communications has constructed Operation Virtual Shield, a network of cameras connected to an operations center. London's ''ring of steel'' incorporates

thousands of cameras, in addition to traffic barriers and inspection points. Boston has hundreds of public and private cameras in its financial district alone.

In 2005, the CTB began designing and implementing a centrally networked municipal camera system. The LMSCC integrates feeds from approximately 6,500 cameras (as of August 2014) that observe public spaces throughout the city. Some of these cameras are owned and maintained by the NYPD, and some by private-sector stakeholders. Of the NYPD cameras, some are fixed (connected via a dedicated fiber optic network) and some are mobile (connected via wireless).

Most cities' cameras are not networked, meaning that an officer must travel to the camera's location after an incident has occurred to retrieve recorded video. The NYPD's cameras, by contrast, are centrally networked so that they are viewable in real time by the officers staffing the LMSCC. Historical video, subject to the restrictions described later, is also easily accessible because the data from the cameras is stored in a central location.

In addition to the raw video, metadata is extracted from many of the cameras in the NYPD's network. Metadata is traditionally defined as ''data about data.'' In the context of video, metadata involves the reduction of the video's content to a more structured form. Video metadata includes tracking information for each object that moves through a camera's field of view, including the object's path, speed, size, and color. This metadata is rapidly searchable via a custom interface and results are presented along with a key frame of the detected object (Hampapur et al., 2007).

License Plate Readers. License plate readers (LPRs) are specialized cameras that record images of motor-vehicle license plates. LPRs are typically coupled with algorithms that extract the alphanumeric characters on the plate from the image; for more details on these algorithms see Anagnostopoulos et al. (2006). The text of each license plate can then be analyzed without the need for a person to examine each plate individually.

The NYPD deploys two types of LPRs: fixed and mobile. The fixed LPRs are being deployed in every lane of traffic at all of the 19 bridges and tunnels that serve as entrance and exit points to Manhattan. They also cover every intersection on Canal Street. The mobile LPRs are mounted on GPS-enabled NYPD vehicles. All LPRs feed pictures of every plate they read in real time back to the LMSCC over the citywide wireless network. The location of the read is also transmitted, whether via the name of the reader (for fixed LPRs) or the latitude and longitude of the reader (for mobile LPRs). As of August 2014, the NYPD receives approximately 2 million reads per day.

Environmental Sensors. The NYPD deploys a variety of environmental sensors throughout the city, including radiation, chemical, biological, and explosives detectors. These sensors are useful for both terrorist attack interdiction and response. The readings from many of the sensors are networked into the LMSCC for increased situational awareness.

As an example of the deployment of environmental sensors, consider the radiation sensors. Similar to the LPRs, the NYPD's radiation sensors come in fixed and mobile varieties. The fixed sensors are deployed throughout the city mounted on the roofs of police facilities. The mobile sensors come in several different types; belt-worn personal radiation detectors as well as more sensitive boat and motor-vehicle mounted sensors are connected over the citywide wireless network. The NYPD also deploys helicopter-mounted radiation sensors that transmit data back to the NYPD's network via broadcast technology. All of the networked radiation sensor readings are viewable in real time by the officer operating the device and by officers in the LMSCC. All radiation sensors report the observed dose rate; a subset of the sensors can also determine the isotope that has emitted the radiation using spectroscopic identification.

Records

The NYPD maintains a variety of electronic databases, including arrests, complaints, 911 calls, parking summonses, and moving violations. The LMSCC obtained electronic copies of many of the NYPD's databases and aggregated the records into a single database. As of

August 2014, this database contains approximately 1.5 billion entries, including documents from both records and sensors. The LMSCC also receives periodic updates of each of the constituent databases as existing records are edited and new records are created. The frequency of updates differs for each database, depending on the information technology policies of the unit that maintains it. For example, the 911 call database is queried every three seconds for new records. In contrast, the complaint records are updated three times per day in a download of the changes to the previous version of the database.

The LMSCC database is maintained in two separate physical locations for geographic redundancy. In normal operations traffic is load balanced across the two sites, but if one site becomes inaccessible the other site can handle the full traffic load. This failover mechanism is also how the NYPD handles hardware failures in the database or the networking equipment.

Each type of record has its own unique combination of fields, though some fields are universal to every record. For example, every type of record has a field denoting the time and date that it was last updated. In contrast, only 911 call records have a field corresponding to the type of an emergency call received. Additionally, records typically have associated geographic locations (usually translated from a street address to a latitude and longitude pair). For example, complaint records have locations corresponding to where the crime took place and arrest records have locations corresponding to where the person was apprehended.

ANALYTIC PROCESSES

The amount of information flowing into LMSCC via the data sources described in the previous section presents a ''big data'' challenge for the NYPD. How can we identify which pieces of information require follow-up and further investigation by officers? How can an investigator effectively search all of the available information? What value does having access to all of this information add?

Analytic processes help the LMSCC manage the incoming data and provide information to the officers in the field. In this section, we will describe three of our processes in detail, so they can be replicated by other police agencies and operations centers.

At the LMSCC, the data sources described previously can be accessed via a custom-built piece of software known as the Domain Awareness System (DAS). The functionality we describe in this section has been fully integrated into the DAS. Note that the DAS is only used to view NYPD records; there is a separate electronic system for entering and altering them. The DAS is the subject of a unique revenue sharing agreement between Microsoft and the NYPD.[b]

Alerting

The first analytic process we will describe is the LMSCC's alerting process, which consists of three steps:

1. Automated alarm
2. Adjudication
3. Action

We will discuss each step in turn.

The first step in the LMSCC's alerting process is an automated alarm. It is a simple computational task to monitor a flow of information and alarm on any piece of data that meets some predefined criteria. The NYPD applies different alarm criteria to its sensors and records. The sensor automated alarm criteria are applied to the readings from a sensor as they come into the LMSCC in real time, whereas the record automated alarm criteria are applied when a new record arrives at the LMSCC, as well as when existing records at the LMSCC are updated.

The NYPD has configured a variety of automated alarms on its collection of sensors. For example, some of the radiation sensors are set to alarm any time they measure an amount of radiation larger than a particular threshold dose rate. Some radiation sensors alarm anytime they identify a particular dangerous isotope, regardless of dose rate. Categorical data, such as that produced by LPRs, requires a slightly different approach. Every LPR read is compared to a watch list of plates; any match triggers an automated alarm. Data from the 911 feed is similarly monitored; any 911 call of specific types

judged to be potentially connected to terrorism (e.g., suspicious packages, hazardous materials, or shots fired in the proximity of a high-risk building) triggers an automated alarm.

Automated alarms can also be applied to camera data sources. The LMSCC has configured automated video analytic alarms to trigger when certain conditions in a camera's field of view are met. For example, an automated alarm will trigger when an object with certain characteristics of a package appears in a camera's field of view and is motionless for a set period of time. Technical details of these abandoned object alarms are available in Bayona et al. (2009). Other configurable video analytic alarms include movement in a frozen zone, crossing of a threshold, and travel against the flow of traffic.

The second step in the alerting process, adjudication, consists of a uniformed officer stationed at the LMSCC closely inspecting the automated alarm. For some sensors, the officers inspect the alert to determine whether it is a true or false positive, and for another set of sensors the officers use their judgment to ascertain whether the alert is potentially connected to terrorism. Some alarms require both determinations to be made. The purpose of this close inspection is to limit the subset of alarms for which follow-up action may need to be taken.

For example, LPR alerts simply require the officers to determine whether the alert is a true or false positive. This is necessary because the plate-reading algorithms mentioned earlier that perform a first cut at extracting the characters from each imaged plate are typically not as accurate as a person inspecting an image. Therefore, the adjudication step consists of an officer confirming that the plate was correctly read by the software. In contrast, 911 alerts are adjudicated primarily to determine whether they are potentially connected to terrorism. For example, an officer will examine the context around a shots fired alert to determine whether it bears the hallmarks of a terrorist attack. Finally, video analytic alerts require an officer to make both sets of judgments. First, the officer inspects the video to determine if the alert is a true positive, because the algorithms may not be as accurate as an officer inspecting the video. For example, a shadow moving across the camera's field of view can be mistakenly interpreted by the video analytic software as an abandoned object. Next, the officer looks at the context in which the alert was triggered to determine if it is potentially connected to terrorism. For example, in the case of an abandoned object, the officer will examine the context around which the object appeared in the camera's field of view.

The third step in the LMSCC's alerting process, action, takes place if, as a result of adjudication, the officer determines that follow-up action is required. The actions NYPD takes based on the alert are prescribed by a carefully crafted concept of operations (CONOPS) tailored to the type of alert requiring a response. For example, a radiation alert adjudicated to be a true positive may prompt the deployment of additional units with specialized equipment. As another example, if an abandoned package is detected by the video analytics software, uniformed officers in the vicinity may be contacted to secure the area and request specialized explosive expertise if necessary.

Investigative Support

The next analytic process we will describe is for information gathering in support of investigations. The CTB is currently in the midst of deploying this capability to the entire NYPD; this will be discussed further later.

Traditionally, as described earlier, law enforcement data has been stored in separate databases organized by source. The independent interfaces make it more difficult for investigators to identify connections across the data sources. The NYPD recognized the inefficiencies inherent in having independent parallel databases and, to eliminate these inefficiencies, built an integrated search engine to search all of the available records and historical sensor information simultaneously. This integrated search is built on an indexed database that returns all of the results within a few seconds. Full Boolean logic for search terms has also been implemented.

To aid the interpretation of the integrated search results, the NYPD has found it helpful to give the investigator the option to view the results with three different data visualizations. The default view is simply to output the results in list form sorted by the date of the record, in

reverse chronological order. Photo view is the first alternative; it displays any photograph associated with the result in a small preview, allowing for a quick visual inspection of LPR images and mugshots from arrest records. Calendar view is the second alternative; it arranges the results into a format similar to an appointment book. This facilitates the quick identification of patterns specific to time of day or day of the week. All of these data visualizations are accompanied by a map displaying the locations of the results.

The integration of historical sensor data and records is particularly useful when supporting investigations. For example, if the whereabouts of a particular vehicle are of interest for a case, it is very illuminating to display in calendar form the LPR reads, parking summons, moving violations, and complaint records in which the vehicle is named.

The construction of a single integrated database also allows the user to quickly review associated records and sensor data, a process the NYPD calls ''correlation.'' For example, if an investigator has identified a license plate read of interest, the database will identify any records with the same license plate in any field. These correlated records may include other LPR reads, DMV records, 911 calls in which the plate is mentioned in the narrative, parking summonses in which that vehicle was cited, and moving violations in which that vehicle was being driven. The database also displays all license plate reads from the same reader that occurred within a window of time around the license plate read of interest. The quick identification of these correlated records, without any additional work by the investigator, improves investigative efficiency.

The integrated search capability has also merged crime mapping and investigative support. The NYPD has had a robust crime mapping capability since the introduction of the CompStat process (discussed previously) in order to identify hotspots and temporal trends. However, the crime mapping software is not designed to allow an investigator to dig deeper into cases; investigators must swap from the crime mapping software to a record database to peruse the records underlying a crime hotspot. Now, the construction of the index that supports integrated search allows an investigator to simultaneously view a crime map (e.g., of arrests or complaints) and the individual records underlying the map without having to access a separate database.

Automated Pattern Recognition

The final analytic process we will discuss is automated pattern recognition for the driving behavior of vehicles on the NYPD's watch list. The raw data for automated pattern recognition are the true positive reads of plates from the watch list, described previously. Automated pattern recognition is not applied to any plates not on the watch list.

The NYPD has developed an analytic process wherein simple automated pattern detection algorithms look through the data for patterns of two types: time-place and routing. The algorithms are executed every morning via a regularly scheduled computer script. Officers then decide how to act on the identified patterns based on the progress of investigations, the availability of resources, and other police objectives.

Time-place patterns are the simplest patterns imaginable. Two examples of time-place patterns are a vehicle that traverses the Brooklyn Bridge into Manhattan every weekday at 0900 HRS and a vehicle that passes an LPR on the Holland Tunnel outbound every Tuesday at 2100 HRS. These patterns are useful because they allow officers to prepare in advance to perform interdictions at the anticipated time or place.

Routing patterns, on the other hand, occur when a vehicle passes multiple LPRs on a repeated trip, even if the trips themselves occur at irregular intervals. For example, consider a vehicle that has been observed on several occasions traversing the Brooklyn Bridge into Manhattan and then 20 minutes later crossing the Holland Tunnel into New Jersey. Even if the time of these trips cannot be predicted, once the LMSCC alerts on the first read of the plate on the Brooklyn Bridge, there is enough warning to deploy officers to interdict the vehicle at the Holland Tunnel 20 minutes later.

Supplementing the identified time-place and routing patterns is a list of license plates on the watch list that have been detected recently, but for which the software was unable

to identify a pattern. If an officer judges one of these plates to be of high priority, he or she can search through the reads of that plate manually to identify a pattern that the software may have missed.

EXAMPLES

The LMSCC has been an integral component of numerous NYPD operations and investigations. In this section, we will describe two examples that demonstrate the value the LMSCC provides to the NYPD and, by proxy, to the public.

Shooting Outside the Empire State Building

At approximately 0900 HRS on August 24, 2012, Jeffrey T. Johnson shot and killed his former coworker on the street outside the Empire State Building in Midtown Manhattan. While fleeing the scene, the gunman confronted two NYPD officers with a loaded weapon; the officers then shot and killed him. Nine additional civilians were wounded.

In the immediate aftermath of a shooting such as this one, conflicting reports make it difficult to establish the facts of what occurred. In this case, early media reports stated that there were multiple shooters and victims, and details regarding the nature of the confrontation conflicted. Many 911 callers also suggested the possibility of an ongoing attack. Reports on traditional and social media were similarly contradictory and confused.

The analytic processes at the LMSCC helped the NYPD quickly determine what had actually happened. At the LMSCC, the many 911 calls regarding shots fired met the predetermined criteria to generate automated alarms from the 911 feed. These alerts were examined by officers in the LMSCC and, given the iconic nature of the calls' location and the previous history of terrorism there, judged to be possibly connected to terrorism. Additionally, the officers tuned the police radio in the LMSCC to the channel corresponding to the 14th precinct where the shooting occurred so they could listen to the first responders on scene.

Officers immediately brought up cameras in the vicinity of the Empire State Building and were able to establish that there was no ongoing attack. This information was provided directly to senior NYPD leadership. The officers then used camera feeds synchronized to just before the 911 calls were made to conduct a camera canvas of the area around the Empire State Building and were able to find video of the officers subduing the gunman. This video was also provided to NYPD leadership, who eventually released it to the media to allay confusion regarding the event.

LPR-Aided Arrest of Suspect with Active Warrant

In policing, a warrant is issued for a suspect when there is probable cause for a law enforcement officer to place that suspect under arrest. Some warrants include additional information to aid officers in finding and arresting a suspect, such as the license plate of a vehicle associated with that suspect.

The NYPD received a warrant for a person suspected of committing wire fraud; this warrant included the license plate of a vehicle owned by that suspect. At the LMSCC, officers entered that plate into the watch list so that it would generate an alert when it was observed by an LPR. The plate was accurately read on several occasions entering Manhattan via the Brooklyn Bridge. A simple pattern analysis demonstrated that the vehicle had on several occasions made that crossing at a specific time and day of the week.

The LMSCC provided this information to field officers, who dispatched a police car to the Brooklyn Bridge at that day and time. The officers in the car encountered the suspect's vehicle as predicted, observed the vehicle commit a violation of traffic laws, and then pulled the vehicle over. The driver of the vehicle was then positively identified as the suspect named in the warrant; he was then placed under arrest.

DISCUSSION

Three aspects of these processes merit further discussion: the implications for information sharing, privacy, and organizational change.

Information Sharing

It is a management cliché that it is challenging to get a large organization to properly communicate among its constituent parts. History has shown this to be true for organizations whose mission includes counterterrorism; for example, at the US federal level, the 9/11 Commission pointed out several failures of information sharing in the intelligence community prior to the attacks and made several recommendations to make similar failures less likely to occur in the future (National Commission on Terrorist Attacks Upon the United States, 2004).

At the NYPD, we have also taken the importance of information sharing to heart. We recognize the necessity of integrating our sensor feeds at a single location; for example, the operational significance of an LPR reading a watch list plate concurrent with a radiation sensor detecting a dangerous source can be more easily appreciated if both alerts are routed to the same personnel. Furthermore, effective information sharing is essential not just for counterterrorism, but also for traditional policing. For example, integrating the data from parking summonses, vehicle registration records, and LPRs is of great use when investigating any crime involving motor vehicles.

The analytic processes we have mentioned in this article aid in information sharing amongst the NYPD's constituent parts. In particular, the process designed to support investigations discussed earlier breaks down information-sharing barriers by making sensor and record information available across the NYPD. The alerting process also discussed previously aids in information sharing by making it easy for everyone to identify the sensor readings of highest import; members of the NYPD do not need to spend hours and hours poring through raw sensor data to find what they need.

Privacy

Concerns regarding counterterrorism programs' impacts on privacy are longstanding (see National Research Council (2008) for further discussion). In 2009, the NYPD developed the Public Security Privacy Guidelines to establish policies and procedures to prescribe the authorized use of the technologies described in this article and to provide for limited access to and proper disposition of stored data (NYPD, 2009b). These guidelines were first released in draft form to allow for public comment.

The guidelines delineate different rules for storing data depending on its source. For example, environmental sensor data, since it has no privacy implications, is preserved indefinitely. On the opposite end of the spectrum, video recorded from the camera network described previously is preserved for 30 days.

Additionally, the privacy policy describes the procedures that the NYPD follows in the event that data must be "archived" beyond the typical timeframe. These procedures involve approvals on a case-by-case basis, depending on the rationale for why archiving was requested. Requests for archiving generally occur when the data is relevant to an investigation. For example, if a robbery is observed on one of the cameras networked into the LMSCC, the NYPD must archive the footage longer than the 30 days specified so that it can be used as evidence during trial.

The DAS software also includes integrated auditing functionality, which allows supervisory officers to ensure that the information contained within the system is only used for legitimate law enforcement and public safety purposes. The NYPD determined that the primary auditing should be performed by an integrity control officer located in the command rather than via a centralized authority because those working side by side have better knowledge of the details of an officer's caseload. There are also centralized units responsible for detecting and preventing misuse of NYPD computer systems in the Information Technology and Internal Affairs Bureaus.

Organizational Change

Taking full advantage of analytics requires modification to long-standing police practices and processes. Most police departments in the United States are organized in a paramilitary structure with rigid chains of command and areas of responsibility. Information in police departments has traditionally been stove-piped, making it difficult for it to be communicated from one part of the organization to another.

The NYPD is in the process of deploying the DAS, including access to the records and sensors described in this article, to all of its precincts and detective squads. This represents a change from the previous analytic model, in which officers had access to a specialized center of excellence known as the real-time crime center (NYPD, 2010). Allowing all officers to have access to these tools is a recognition that everyone is an analyst, not just the officers in a specialized analytic unit.

CONCLUSION

The analytic processes described in this article represent a substantial augmentation of the NYPD's capabilities. It was necessary to develop these processes to make effective use of the large volume of data the NYPD ingests. Because it is the largest police agency in the United States located in an iconic city, the NYPD often encounters challenges in advance of other US and international police agencies. The lessons the NYPD learns are often relevant to these other agencies in other cities. By sharing these lessons in this article, we hope that the other agencies will benefit from our work using analytics to take advantage of data.

ACKNOWLEDGMENTS

The authors would like to thank Police Commissioner William J. Bratton, Deputy Commissioner John J. Miller, Chief James Waters, Richard Schroeder, Deputy Chief Salvatore DiPace, Courtney MacGregor, Lauralee Giovanella, Deputy Inspector Anthony Tasso, Captain Christine Doherty, Captain Brendan Deery, Lieutenant Michael Joy, Sergeant Gordon Pekusic, Detective Dan Higgins, Detective Todd Metro, and Daniel Kusrow, as well as Richard Falkenrath (ret.), Richard Daddario (ret.), Deputy Inspector Andrew Savino (ret.), Nathaniel Young (ret.), Sarah Watson (ret.), and Ryan Merola (ret.). The work described in this article was completed with funding provided by the U.S. Department of Homeland Security's Homeland Security Grant Program.

NOTES

[a] The LMSCC also facilitates partnership between the public and private sectors by co-locating security officers from public and private sector organizations.

[b] For more information on the DAS and the revenue sharing agreement, see the official statement (NYC Mayor's Office, 2012).

REFERENCES

Anagnostopoulos, C. N. E., Anagnostopoulos, I. E., Loumos, V., and Kayafas, E. 2006. A License Plate Recognition Algorithm for Intelligent Transportation System Applications, *IEEE Transactions on Intelligent Transportation Systems*, Vol 7, No 3, 377–392.

Bayona, A., SanMiguel, J. C., and Martinez, J. M. 2009. Comparative Evaluation of Stationary Foreground Object Detection Algorithms Based on Background Subtraction Techniques, *Proceedings of the 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*, IEEE, 25–30.

Beers, R. 2013. Congressional Testimony, House Committee on Energy and Commerce, Subcommittee on Environment and the Economy.

Davenport, T. H., and Harris, J. G. 2007. Competing on Analytics. Harvard Business School, Cambridge, Massachusetts.

Davenport, T. H., Harris, J. G., and Morison, R. 2010. Analytics at Work. Harvard Business School, Cambridge, Massachusetts.

FEMA. 2014. FY 2014 Homeland Security Grant Program. URL http://www.fema.gov/fy-2014-homeland-security-grant-program-hsgp.

Hampapur, A., Brown, L., Feris, R., Senior, A., Shu, C., Tian, Y., Zhai, Y., and Lu, M. 2007. Searching Surveillance Video, *Proceedings of the IEEE Conference on Advanced Video and Signal Based Surveillance*, IEEE, 75–80.

Henry, V. E. 2002. *The CompStat Paradigm.* Looseleaf Law Publications, Flushing, New York.

Langworthy, R. H., ed. 1999. *Measuring What Matters: Proceedings from the Policing Research Institute Meetings.* National Institute of Justice.

National Commission on Terrorist Attacks Upon the United States. 2004. *The 9/11 Commission Report.* US Government Printing Office.

National Research Council. 2004. *Fairness and Effectiveness in Policing: The Evidence.* National Academies Press, Washington, DC.

National Research Council. 2008. *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Assessment.* National Academies Press, Washington, DC.

New York City. 2012. Adopted 2013 Financial Plan.

NYC Mayor's Office. 2012. Mayor Bloomberg, Police Commissioner Kelly and Microsoft Unveil New, State-of-the-Art Law Enforcement Technology that Aggregates and Analyzes Existing Public Safety Data in Real Time to Provide a Comprehensive View of Potential Threats and Criminal Activity. Press release, August 8.

NYPD. 2009a. Engineering Security: Protective Design for High Risk Buildings.

NYPD. 2009b. Public Security Privacy Guidelines.

NYPD. 2010. Best Practice—Real Time Crime Center—Centralized Crime Data System.

Reaves, B. A. 2010. Local Police Departments, 2007. US Department of Justice, Washington, DC.

Silverman, E. 1999. *NYPD Battles Crime: Innovative Strategies in Policing.* Northeastern, Boston, Massachusetts.

Weisburd, D., Mastrofski, S. D., McNally, A. M., Greenspan, R., and Willis, J. J. 2003. Reforming to Preserve: CompStat and Strategic Problem Solving in American Policing, *Criminology and Public Policy*, Vol 2, No 3, 421–456.

Zimring, F. E. 2012. *The City That Became Safe.* Oxford University Press, Oxford, England.