Nayar Prize II
Year 2 Progress Report

A Data-Driven Crime Prevention Program

Principal investigators:

Miles N. Wernick
Motorola Endowed Chair Professor of Electrical Engineering
Professor of Biomedical Engineering
Director, Medical Imaging Research Center

Lori Andrews
University Distinguished Professor
Professor of Law
Director of the Institute for Science, Law and Technology

Yongyi Yang
Perlstein Endowed Chair Professor of Electrical Engineering
Professor of Biomedical Engineering

Student participants:

David Haro Alonso
Ph.D. candidate, Department of Electrical and Computer Engineering; degree awarded in May 2018.

Pauline Panayi
J.D. candidate, Chicago-Kent College of Law

Harshini Tippareddy
J.D. candidate, Chicago-Kent College of Law

Raven Zeer
J.D. candidate, Chicago-Kent College of Law; degree awarded in May 2018.

Elgin Police Department:

Bill Wolf
Deputy Chief

JoAnn Stingley
Social Services Supervisor

Pastor Robert Whitt
Community Liaison

September 27, 2018

Table of Contents

Project Overview

The objective of "A Data-Driven Crime Prevention Program" was to design, implement, and deploy a flexible, innovative model for crime prevention that can be translated to a wide array of communities in the U.S. and beyond, thereby achieving far-reaching societal impact. The project has succeeded in producing advances in predictive modeling technology based on machine learning, along with a legal-ethical framework for appropriately employing this technology to crime prevention in a way that respects constitutional rights and achieves acceptance by the community. The project was undertaken with the Elgin (Illinois) Police Department, which we selected because of its exceptionally progressive culture, positive relationship with its community, and willingness to devote substantial effort to work closely with us in a genuine partnership.

Our project never would have occurred were it not for the Nayar Prize competition. The Nayar Prize funding inspired us to bring together faculty and students from biomedical engineering, electrical and computer engineering, and law and to create a new approach to crime prevention drawing upon the disparate skill sets of the participants. In related work, Professors Miles Wernick and Yongyi Yang had previously developed numerous innovative machine learning approaches to timely problems in the field of medicine. Professor Lori Andrews had pioneered analyses of emerging technologies, ranging from genetic technologies and nanotechnologies to digital technologies. She had previously explored legal and ethical constraints on crime prediction based on alleged genetic predispositions to commit crimes. The project was a true collaboration across disciplines; the participants from the legal field made suggestions about which variables to use in the algorithm and how to weight them, while the participants from engineering were involved in the creation of the policy guidelines for the Elgin Police Department and in the development of legal and ethical guidelines for judges, lawyers, for-profit companies, and university researchers to use when implementing predictive algorithms in the criminal justice setting. Researchers from the various disciplines spoke about the project to scientific, legal, public, and media audiences and also presented information about the project to the Illinois Tech community.

Our crime-prevention program parallels current directions in the field of preventive medicine, where it has been recognized that intervention strategies should shift from a reactive stance to one that is proactive. Intervention strategies in preventive medicine often begin with statistical models that estimate an individual's risk for developing disease based on various risk factors (such as smoking or a genetic predisposition to lung cancer), and then deploy treatment (such as smoking cessation programs) to reduce future risk.

Some of the ideas developed in the project were inspired by the team's work in the medical field, and vice versa. During this project, our student David Haro Alonso applied ideas developed in this project to an important risk prediction problem in medicine.[1] An editorial[2] was recently published in a prominent medical journal that was entirely devoted to praising that work and showing how it presents keys to applying machine learning to risk prediction.

In an approach analogous to preventive medicine, we are developing risk models, founded on machine learning technology, that assess an individual's future risk of becoming a victim or arrestee in a crime incident. In our intervention strategy, individuals identified by the model to be at high risk are invited to work with social workers to help identify which social services programs would be most helpful to them (these may include job training, drug rehabilitation, and educational opportunities). Social workers then follow up with the individual consistently to ensure success. As in preventive medicine, the aim is to reduce the individual's future risk by modifying the circumstances that gave rise to that risk.

The overall goal of our project is to help people to better their lives, thereby benefiting them and their families, and ultimately leading to a safer community. Our objective is to demonstrate an approach to crime prevention that will serve as a model to other communities nationwide and throughout the world. We will disseminate knowledge gained in the project through scientific publications; law review articles; training for police officers, community leaders, and judges; media interviews and/or op-ed pieces; and interactions with others interested in using risk modeling in their social services or criminal justice activities.

Executive Summary of Progress Report

We began by assessing the needs of the Elgin community, working with a team of representatives of the Elgin Police Department (EPD), including the deputy chief of police and other subject-matter experts, the department's head social worker, and a local pastor, who acts as a liaison to the community and local churches. We discussed and selected strategies in a collaborative and iterative process involving our legal and technical experts, our students, and the EPD. We conducted a detailed legal and ethical analysis of relevant case law to guide the development of the risk model and the intervention. We undertook technical development of several preliminary versions of the machine learning risk model and, after extensive consultation with the EPD, we selected one of the versions to become the basis for the first pilot phase of interventions. We performed a thorough statistical analysis of the risk model, not only to test its accuracy, but also to ensure that no racial or ethnic bias was present in the results. We also evaluated the risk model by computing risk assessments for individuals in the past to see whether experts from within the EPD would agree that the assessments were reasonable and insightful. This step was a resounding success, with experts reporting that the model produced recommendations that pointed to individuals who later were involved in very serious crimes, or who had been successfully assisted by previous social service outreach programs. In a fully collaborative effort, the EPD and our team developed a Standard Operating Procedure (SOP) document to define the goals and procedures for implementation of the intervention program. The SOP document received final approval by the City of Elgin. We have thoroughly reviewed the intervention concept and SOP with the head social worker and pastor, and the logistics for making decisions about particular individuals were decided. The pastor shared our program's plans at a meeting of local church leaders to obtain feedback and community buy-in. His deep familiarity with the local community and his ability to talk to at-risk individuals in a way that resonates with them made him an invaluable resource to engage them in seeking the help that they need. In addition to our specific work with Elgin, we began disseminating knowledge gained and joining the national dialogue on crime. We participated in media interviews relating to crime prevention (e.g., *60 Minutes* and *The New York Times*), we gave invited talks to classes at Illinois Tech and other universities and at national and international meetings, we initiated discussion with other interested institutions and programs, and we began preparation of academic publications based on our work.

The following sections provide details of our progress so far and our plans for Year 3 if we were selected for funding.

Collaboration with the Elgin Police Department (EPD)

Background

In the law-enforcement community, new approaches become nationwide "best practices" after being successfully demonstrated by a city; therefore, it is important to choose that first city wisely so as to create the best precedent possible. Of course, a positive outcome in a pilot study is the ideal basis for future applications by social services organizations as well. Thus, we chose to work with the Elgin Police Department (rather than that of a different city) as the initial proving ground because of its history of innovative, culturally-sensitive, community-based programs to prevent crime and create a safer, more productive community. Our selection proved to be even more fortuitous than expected, with the EPD fully incorporating our recommendations at every step of the way and showing great sensitivity and progressive thinking.

In the late 1980s, Elgin's crime problem was so severe that two documentaries were made about it. In response, the EPD adopted several strategies to prevent crimes. For example, EPD began an officer-in-residence program in which police officers were given housing in the most dangerous parts of the city. Initially, the program posed great risk to the officers involved (their houses were set afire and their cars blown up), but ultimately, the individual police officers became integrated into their local communities. Seeing first-hand what was happening on the streets and in the schools, they were able to help develop programs to serve the needs of community members. The previously dangerous neighborhoods that officers volunteered to live in have now become much safer. But Elgin faces other problems—crimes by the homeless and those who are mentally ill, gang problems (with all the gangs in Chicago having members in Elgin as well), and other serious crimes.

For the past 27 years, the EPD has maintained a social work division that runs programs in the high schools, matches parolees with social services programs that can benefit them (such as drug rehabilitation and GED programs), and, in domestic violence cases, provides extensive aid attuned to the specific needs of the individuals, such as when the mortgage crisis triggered increased domestic violence.

In working with the EPD, we have been impressed by their commitment to provide relevant social services—*with consistent follow-through*—to people who are identified as being at risk. In addition, the EPD worked with us to develop a Standard Operating Procedure titled "Social Services Crime Prevention Program Procedures." This SOP meets the legal and ethical guidelines that we developed through our legal-ethical analyses, and provides clear rationales and processes for the administration of the program, its interaction with participants, its transparency, and its obligations to the community (see Appendix A).

Understanding Elgin's Needs and Resources

The EPD's motivation to initiate the project arises from the recognition that a relatively small fraction of Elgin's population accounts for most of the serious crime activity; thus, by helping even a few people to reduce their involvement in crime, the EPD feels that the program can have a significant impact on the community.

The EPD explained to us that Elgin's crime problems are more typical of small- to medium-sized U.S. cities, such as those found in parts of the country where the local economy is not as vibrant as it once was. The EPD's deputy chief and its subject-matter experts discussed the project goals and proposed a list of crimes that they consider to be serious (see Table 1).

The EPD feels that our approach helps them to identify individuals in a thorough, systematic, and data-driven way, whereas currently only a subset of at-risk individuals are referred for outreach (such as when police or social workers happen to come across these individuals through their daily activities). While their current outreach efforts are successful, the EPD feels that it is important to reach out to individuals who are at extreme risk, but whose situations may be overlooked because they are simply unknown to the social workers.

An individual may have a spectrum of needs based on his or her circumstances, with assistance depending in part on whether the individual is a victim or arrestee. However, we decided that the judgment as to what assistance would be most appropriate should be made on a case-by-case basis through a careful discussion between the individual and a social worker. Therefore, we have decided that the risk model would not seek to differentiate victims from arrestees. Such an approach also makes sense in view of the fact that those at risk for involvement in serious crimes often become both victims *and* perpetrators within a given time period.

The EPD has offered substantial personnel resources to the project. Deputy Chief Bill Wolf, our initial point of contact for the project, was responsible for top-level policy making and discussions about the project direction. The EPD also involved several high-ranking police officers, who have first-hand interactions with at-risk individuals and who have expert knowledge of the neighborhoods and challenges. The EPD also assigned a very capable member of their information technology team who has assisted with access to relevant crime records.

On a day-to-day basis, the most important EPD personnel in the program have been a team of social workers, led by JoAnn Stingley, Social Services Supervisor, and Pastor Robert Whitt, Community Liaison. Ms. Stingley, the head social worker of the EPD, has been applauded in the media as an "unsung hero"[3] for her selfless community activities, including leading fundraisers that support many of the social programs in Elgin. Pastor Whitt, a member of Elgin's Human Rights Commission, is a well-known community leader who engages in efforts to promote healthy communications between police and the public and to foster racial harmony, such as through the community meeting he organized to help people of various backgrounds to understand the Black Lives Matter movement.[4]

We have found the EPD team to be very impressive, showing a remarkable degree of sensitivity that is in stark contrast to the negative public image of some police departments. The EPD places a great deal of emphasis on community-oriented policing, social work, and addressing the underlying needs of people involved in crime, rather than treating policing as a process of simply arresting people. Our conversations with various Elgin police officers have demonstrated that this progressive culture extends down to the rank-and-file officers.

Legal and Ethical Research

To provide a sound basis for development of our risk model and the associated social services intervention program, we conducted a thorough legal and ethical analysis relating to the uses of algorithms in relation to crime. A disturbing trend is emerging in which algorithms are being used, or at least considered, in ways that we believe are inappropriate, including decisions we categorize as punitive, such as those affecting arrests, convictions, sentencing, or parole. Any use of an algorithm in the criminal justice system must meet the strictest of legal and ethical standards so as to protect the rights of individuals affected by such decisions.

Thus, to guide the development of our risk model and intervention program, and to educate the public (especially judges, police, and prosecutors), we set out to delineate the issues that must be considered when employing algorithms in the context of law enforcement and the criminal justice system. Toward this goal, we conducted a detailed legal and ethical analysis of U.S. Constitutional principles and of all the appellate court cases (n=101) that reviewed the factors that police officers said they relied on to make predictions about criminality.

We analyzed the 101 court cases that dealt with law enforcement prediction of criminality recorded the variables used in the predictions and analyzed how courts determined whether it was legally permissible to use each variable. We found a total of 125 variables that law enforcement officers said they used to predict that a particular individual was likely to commit a crime or have committed a crime. In the cases we considered, the number of variables that were employed ranged from one to eleven. The court cases, ranging from the trial court level to U.S. Supreme Court, emphasized the importance of context, suggesting that the use of multiple variables at once is the most appropriate approach. For example, while a certain single variable used alone (such as location) is deemed inappropriate by courts, it can be permissible when used with other predictive variables.

We found that courts were more lenient toward law enforcement officials with respect to the factors they used to predict potential dangerousness at airports than elsewhere in public. We determined that the precedents for the public, non-airport settings were more appropriate for our use. (See Appendix B). Those cases emphasize that it is legally improper to use race as a predictive variable and we have thus designed our risk models in a way that does not use race as a variable, nor does it use factors that are linked closely to race, such as place of residence or gang affiliation.

Even beyond its usefulness within this project, our extensive analysis is sufficiently important and unique that we will expand upon it and publish it as an article to guide judges in their work and to guide police departments in their training of officers.

Recommended Legal and Ethical Guidelines

Based on our analysis of cases and Constitutional principles, we fashioned guidelines for the use of predictive algorithms and we developed our risk model and intervention strategy so as to meet these legal and ethical guidelines. These guidelines are summarized as follows:

Algorithms should not be used by the criminal justice system in punitive ways, such as for sentencing or parole decisions. Crime-risk models should be designed to help people to avoid adverse outcomes, not as a basis to punish them. Constitutional due process requires that decisions that punish an individual should be based instead on analyses that are entirely specific to the individual,[5] and not on aggregate estimates of risk in relation to historical data. Algorithms that predict recidivism are inappropriate for sentencing decisions because they focus on the wrong scientific questions. The appropriate question is whether a longer sentence is more likely to deter future crime in a particular individual. That may depend on factors such as mental illness, or whether incarceration itself leads to a greater risk of criminality (for example, by providing opportunities for the individual to learn additional criminal skills). Most importantly, the Constitutional right of due process requires that the defendant receive an "individualized" sentence and not be sentenced based on aggregate risk predictions. In the 2016 *State v. Loomis* case, the court ruled that "using a risk assessment tool to determine the length and severity of a sentence is a poor fit. As scholars have observed, '[a]ssessing the risk of future crime plays no role in sentencing decisions based solely on backward-looking perceptions of blameworthiness, … is not relevant to deterrence, … and should not be used to sentence offenders to more time than they morally deserve.'"[6] We will be using our risk assessment tool only in instances that offer social services, not for punitive actions such as increasing sentences.

Algorithms must avoid racial bias. The development of an algorithm can be a way to avoid the idiosyncratic (or even racially-biased) decisions that could be made by police officers in the absence of such a scientific model. Yet care must be taken not to introduce racial biases into the algorithm itself. In Pennsylvania, state police instructed bank employees to take photos of suspicious-looking African-Americans—thus setting the stage for creation of a criminal profile that applied only to African-Americans. This action was held to create a cause of action under the Civil Rights Act on behalf of an African-American man who was photographed.[7] Similarly, courts have held that race cannot be used as a predictive factor in police decisions to stop and frisk an individual. The U.S. Supreme Court has held race to be an impermissible factor when assessing reasonable suspicion,[8] saying that "the Constitution prohibits selective enforcement of the law based on considerations such as race."[9]

In a 2017 U.S. Supreme Court case, *Buck v. Davis*, the defendant's expert argued that race is a factor "know[n] to predict future dangerousness."[10] The Court wrote that "this is a disturbing departure from a basic premise of our criminal justice system: Our law punishes people for what they do, not who they are."[8] We argue that, if carefully designed, a risk model can help to make decision-making fairer, not less fair, by avoiding the idiosyncratic or even racially biased notions that humans may entertain in the absence of a scientific approach based on hard data.

We not only excluded variables relating to race, but we also analyzed the output of our risk model to ensure that racial or ethnic bias did not arise inadvertently, as we explain later.

Algorithms, and procedures for their use in crime prevention programs, should be publicly disclosed. Numerous for-profit companies offer predictive policing services that involve a proprietary algorithm. If decisions are made about individuals based on a secret algorithm, this will infringe their Constitutional rights. In 1977, the U.S. Supreme Court held that a defendant's due process rights had been violated when he was sentenced to death based on a report which contained confidential passages that he was not allowed to see or refute.[11]

In an analogous situation, when a company developed a forensic DNA test and refused to disclose its scientific underpinnings to a criminal defendant, a Vermont court prohibited its use on the grounds that the defendant had a right to know how the predictions were made about DNA matching and a right to challenge the methodology.[12] The court said that such secrecy "is more than problematic, it is anti-scientific in that it inhibits the ability of scientists in the field (including defense experts) to test the manufacturer's claims."[10]

In keeping with this recommendation, we are making the full technical details of our model publicly available, and Pastor Whitt is actively working to inform the community about the details of the program.

The algorithm must not use variables in ways that infringe rights of free speech, association, and privacy. The most appropriate sources of data for crime-risk modeling are prior arrests and victimization. However, there is an increasing interest among police departments in using a person's social media posts to investigate crimes and predict criminality. Seventy percent of law enforcement agencies use social media posts in criminal investigations.[13] In recent years, several companies sold services to police departments that use social media posts to predict an individual's dangerousness[14] or to monitor activist groups, such as black activists.[15]

The use of a variable based on a social media post in a predictive policing model is scientifically unsound[16] and legally improper. Using a variable based on social media posts in a predictive algorithm could violate a person's right of freedom of expression and freedom of association. It could also violate the rights of the friends and relatives who communicate with him on social media. In an analogous situation, agents of the Federal Bureau of Investigation went to court to ask permission to remotely activate a suspect's webcam to take photos of what the suspect looked like and screenshots of what he was doing on his computer. In that case, the request was turned down because of the intrusiveness of the surveillance and the fact that innocent people who used the same computer (which might have been in a home, a dorm, or even an internet café) might have been photographed as well.[17] Similarly, the American Civil Liberties Union successfully convinced Twitter, Instagram, and Facebook to cease providing data to the companies (Geofeedia, SnapTrends, and Media Sonar) offering predictive services based on social media posts to police departments.[13]

We have carefully avoided the use of any kind of information that goes beyond standard crime records so as to avoid the pitfall illustrated by the aforementioned cases.

Development of the Risk Models

The aim of our current risk model is to estimate the overall severity of crime involvement that an individual is likely to experience in the upcoming year, as a victim or arrestee. This risk estimate is based on the individual's involvement in crime incidents in the preceding two years. The risk model is developed and validated using historical data, then applied to current data to compute estimated future risks of individuals. In this section, we describe the process by which the model was developed and present the results of that development.

We have been working closely with the Elgin Police Department (EPD) to predict which people are at highest risk of being involved in a serious crime as a victim or arrestee in order to offer social services to those individuals in an attempt to avert future crime and victimization.

We began by asking EPD to provide a list of crimes that would be considered serious enough to merit inclusion in a risk model. Their list is shown in Table 1. In the following discussion, we detail an iterative development process that led to our current solution.

Table 1. Serious crimes identified by Elgin Police Department

CODE	DESCRIPTION
	Homicide
0110	1st-degree murder
0115	Homicide, unborn child
0120	Voluntary manslaughter, unborn child
0130	2nd-degree murder
0141	Involuntary manslaughter
0142	Reckless homicide
0150	Justifiable homicide
0160	Concealment of a homicide death
0165	Involuntary manslaughter & reckless homicide of an unborn child
0170	Drug induced homicide
0190	Solicitation for murder or murder for hire
	Criminal sexual assault
0260	Criminal sexual assault
0261	Aggravated criminal sexual assault
0262	Forcible sodomy
0280	Predatory criminal sexual assault of a child
0281	Criminal sexual assault with an object
	Robbery
0310	Armed robbery
0320	Robbery
0325	Vehicular hijacking
0326	Aggravated vehicular hijacking
0330	Aggravated robbery
	Battery
0410	Aggravated battery
0460	Battery
0470	Reckless conduct
0475	Battery of an unborn child
0480	Heinous battery
0485	Aggravated battery of a child
0486	Domestic battery
0487	Aggravated battery of an unborn child
0488	Aggravated domestic battery
0491	Aggravated stalking
0495	Aggravated battery of a senior citizen
	Assault
0510	Aggravated assault
0560	Assault

Version 0 (V0)

In initial attempts to construct a risk model, we defined the problem as one of predicting the likelihood that a person with an arrest history in the past two years would be involved as a victim or arrestee in a "serious" crime in the upcoming year. This initial problem statement roughly mirrored that used in a risk model our team had made previously for Chicago, but it encompassed a broader array of "serious" crimes than did the risk model for Chicago, which only involved shootings and homicides. While this approach produced risk assessments that were indeed predictive, they were not sufficiently specific to be useful. That is, a large fraction of recent arrestees will be involved in a future "serious" crime, because the definition of "serious" is so broad. Thus, the model failed to distinguish degrees of risk when recommending persons for the social workers to contact, and it identified far more people than the social workers could handle.

We realized that the problem with our V0 approaches was that they made no distinction between the most and least serious of the "serious" crimes, nor did they take into account the number of crimes in which a person was involved. Thus, for example, a person who was shot twice recently would be treated as having the same crime risk as a person who was arrested just once for disorderly conduct.

Version 1 (V1)

To refine the problem statement, we recast the problem in V1 as one of predicting an overall measure of the severity and amount of crime that a person has recently been involved in, which we term the *crime-severity index*.

To define the crime-severity index, we began by assigning a severity score (*crime-severity value*) for each of the "serious" crime types. We defined the crime-severity value for each crime type to be proportional to the maximum sentence prescribed for that crime in the State of Illinois (see Table 2). The reasoning behind this approach is that maximum sentences reflect our society's notion of the relative severity of different crimes.

Table 2. Crime-severity values

Crime-severity value	Code	Description
100	0110	First degree murder
100	0115	Homicide of an unborn child
60	0280	Predatory criminal sexual assault of a child
30	0170	Drug induced homicide
30	0190	Solicitation for murder or murder for hire
30	0261	Aggravated criminal sexual assault
30	0281	Criminal sexual assault with an object
30	0310	Armed robbery
30	0326	Aggravated vehicular hijacking
30	0480	Heinous battery
30	0485	Aggravated battery of a child
20	0262	Forcible sodomy
15	0120	Voluntary manslaughter of an unborn child
15	0130	Second degree murder
15	0260	Criminal sexual assault
15	0325	Vehicular hijacking
15	0330	Aggravated robbery
7	0320	Robbery
7	0487	Aggravated battery of an unborn child
7	0495	Aggravated battery of a senior citizen
5	0141	Involuntary manslaughter
5	0142	Reckless homicide
5	0160	Concealment of a homicide death
5	0165	Involuntary manslaughter and reckless homicide of an unborn child
5	0410	Aggravated battery
5	0488	Aggravated domestic battery
5	0491	Aggravated stalking
1	0460	Battery
1	0470	Reckless conduct
1	0475	Battery of an unborn child
1	0486	Domestic battery
1	0510	Aggravated assault
0.1	0560	Assault
0	0150	Justifiable homicide

To illustrate how a person's crime-severity index is computed, suppose that the person has been involved as a victim or arrestee in a homicide and two armed robberies. Because the crime-severity values for these two crimes are 100 and 30, respectively, the person's crime-severity index would be $1 \times 100 + 2 \times 30 = 160$. Expressed mathematically, if a person is involved in n_i crimes of type i, $i = 1, \dots, N$, and if a crime of type i has crime-severity value c_i, then the individual's overall crime-severity index is $CSI = \sum_{i=1}^{N} c_i n_i$.

Among individuals with a non-zero crime-risk index, the highest value seen in historical data for any given one-year period is 100. The values corresponding to the 25th, 50th and 75th percentiles are 1, 1, and 2, respectively.

Data and Features

Our analyses were based on a historical database of anonymized crime incidents provided by EPD, which included information on a total of 6,828 crime incidents during the period from January 2012 to July 2016. We collected information from 14,736 cases, of which 755 were involved in future "serious" incidents in the following year. For privacy protection, individuals were identified not by name, but by a unique identification number (which EPD refers to as a *JacketID*), with other sensitive information concerning sex, race, and place of residence being omitted from the data set and the modeling research.

Working with EPD, we identified a set of 20 variables (known in machine learning as *features*) that might be predictive of future crime risk. These features were all based on the crime-incident data provided by EPD.

We introduced three additional features, motivated by our prior research, in which it has been well established that if a person (Person A) is often arrested together (*co-arrested*) with another person (Person B) who is frequently involved in crime, then Person A is at increased crime risk in the future, either as a victim or arrestee. Thus, connections in a person's arrest history are known to give a clue as to future risk.

These connections can be summarized by a *social network* diagram, which depicts the co-arrest connections among the individuals in the data set. An example of such a diagram is shown in Figure 1.

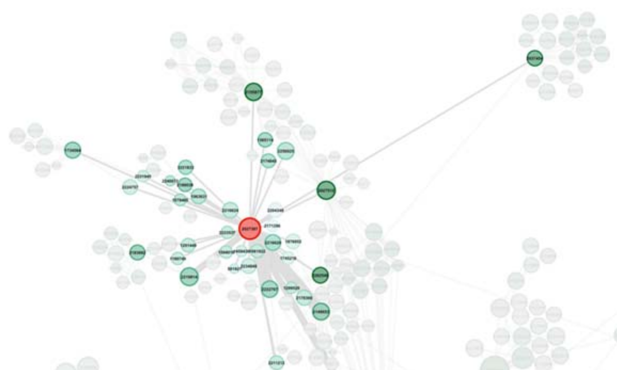

Figure 1. Example of a social network diagram, showing co-arrest connections among at-risk individuals. Highlighted individual (red circle) is at increased risk due to co-arrests with other at-risk persons (green circles).

Mathematically speaking, a social network is a weighted and undirected graph in which the individuals are represented by nodes, and the connections are represented by edges. An individual's connectedness can be described in terms of his or her *geodesic distance* from other arrestees in the social network, described as follows. If Person A and Person B have been co-arrested, they are said to have a first-degree connection (or geodesic distance of 1). Persons who have been co-arrested with Person B are said to have second-degree connections to Person A (or geodesic distance of 2). From the social network, we computed three features: 1) the number of first-degree connections, 2) the number of second-degree connections, and 3) a binary feature specifying whether each individual is a member of the largest interconnected subset of all the individuals in the data set.

Table 3 lists the 20 crime-incident features and three social-network features from which our risk model was permitted to choose a subset that were most predictive. Table 3 also shows the minimum, maximum, mean, and standard deviation of each of the 23 features.

Table 3. List of features considered in the crime-risk modeling process

	Features	Min	Max	Mean ± Std. Dev.
	Features based on individual's crime statistics			
1	Frequency of being victim or arrestee	1	17	1.42 ± 1.00
2	Age at latest arrest	10	85	30.35 ± 12.60
3	Trend in crime incidents	-1.8	2.3	0 ± 0.30
4	Homicide arrests	0	1	0 ± 0.05
5	Sexual assault involvement	0	3	0.04 ± 0.21
6	Robbery involvement	0	3	0.04 ± 0.22
7	Battery involvement	0	17	1.26 ± 1.00
8	Assault involvement	0	3	0.12 ± 0.35
9	Domestic involvement	0	14	0.67 ± 0.85
10	Booking	0	15	1 ± 1.39
11	UUW arrests	0	1	0 ± 0.06
12	Home invasion involvement	0	1	0 ± 0.06
13	Juvenile arrestee	0	1	0 ± 0.04
14	Burglary involvement	0	1	0 ± 0.04
15	Shooter	0	1	0 ±0.02
16	Shot	0	1	0 ± 0.01
17	Violent crime involvement	0	6	0.27 ± 0.52
18	Felony arrests	0	5	0.23 ± 0.46
19	Local arrestee	0	1	0.01 ± 0.08
20	Misdemeanor involvement	0	17	1.26 ± 0.99
	Features based on social network			
21	Number of co-arrestees (1 degree separation)	0	30	0.83 ± 1.93
22	Co-arrestees within 2 degrees of separation	0	322	3.2 ± 0.19
23	Individual is a member of the largest cluster of connected co-arrestees (yes/no)	0	1	0.04 ± 0.19

Model Selection and Optimization

In machine learning, it is not possible to determine *a priori* which form of mathematical model will work best in a given data set, because the data do not obey a probability distribution that can be written analytically, and the dimensionality of the data prevent them from being fully visualized. Thus, it is standard practice to select the best model by using historical data for which the correct answer is known (i.e., the particular person was a victim or arrestee within the next 12 months) to compare the performance of several models using a pre-defined performance metric. Before describing the winning model, let us first introduce the performance metric used to guide the process. This choice is pivotal, because it determines which model will be selected. It also guides the optimization of model performance.

Performance Metrics

The goal of the risk model is to prioritize the social-services outreach program by ranking individuals in descending order of predicted crime-severity index. This is an example of *learning to rank* (or *machine-learned ranking*), an application of machine learning to rank a list of elements.[18] We use two metrics to quantify the success of our model's ranking of risk. First, for use in interpretation by users, we employ *precision* for the 20 highest-ranked persons, i.e., the fraction of these 20 persons who will actually be involved in a serious crime. However, more importantly, for optimizing and testing our model, we measure performance by using the *normalized discounted cumulative gain* (NDCG), a metric that is widely used to measure the effectiveness of search algorithms that rank the relevance of web pages.[19]

For quantitative purposes, NDCG has two advantages over the precision metric. First, it allows individuals to have a continuous-valued relevance or importance (overall crime severity, in our case), while the precision is based only on a binary outcome (involved in a serious crime or not.) Second, NDCG involves a discount function over the rank, which penalizes the appearance of relevant individuals in lower positions in the ranking,[20] whereas the precision uniformly weights all positions. Figure 2 shows an example of the difference between precision and NDCG, and the values that we would obtain in three hypothetical situations.

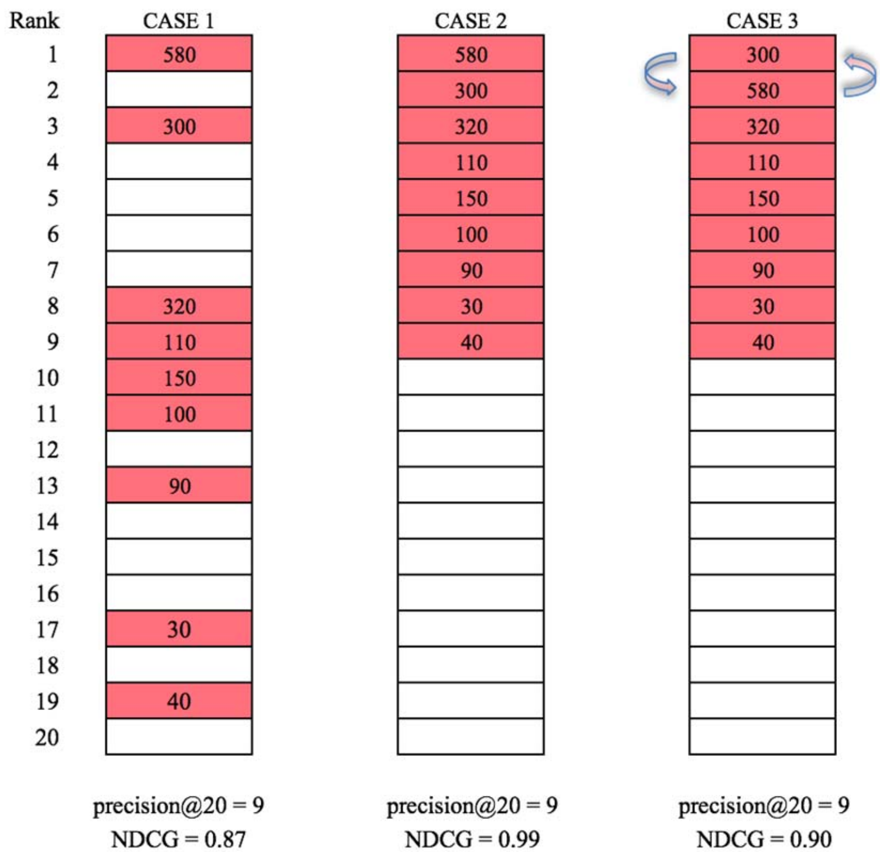

Figure 2. Comparison between precision and NDCG metric. The figure shows that the NDCG is sensitive to ranking performance within the top group, while precision is not. Rankings for the three modeling cases shown have the same precision for the top 20, but different values of the NDCG, which considers the individual's position in the rank and his or her relevance. Cases 2 and 3 are almost the same, but the first two persons' rankings are interchanged in Case 3 to show how the NDCG decreases when the second-highest ranked person is mistakenly ranked first.

The NDCG metric is constrained to values between 0% and 100%, with 100% reflecting the perfect ranking. NDCG is the normalized version of the *discounted cumulative gain* (DCG), which is defined in equation (3.1) as the sum of all the individuals' crime-risk index values, discounted by their ranking, where the discount factor depends on the rank position *i* to which the individual is assigned:

$$DCG = \sum_{i=1}^{N} \frac{crime - risk\ index_i}{log_2(i+1)} \tag{3.1}$$

We employ the widely used logarithmic factor in the NDCG because it produces a smooth reduction of the scores that matches the user's perception of relevancy (see Figure 3). By doing so, the most severe incidents quickly lose importance as their position ranking is lowered.

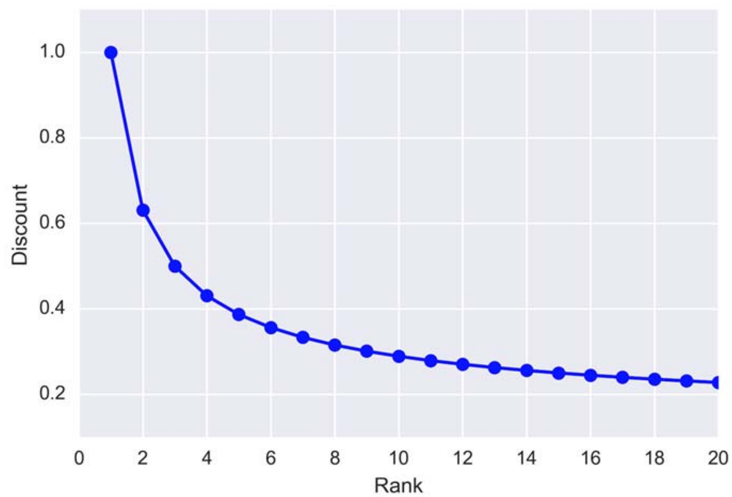

Figure 3. NDCG discount function. The shape of the function shows that the relevance of an individual decreases exponentially with the rank.

Normalizing the DCG allows us to compare the performance of different models. The metric is normalized by the ideal DCG, which is the DCG of the best possible ranking.

Data Normalization

The ranking problem is reduced to a regression algorithm that predicts the individuals' crime-risk index values and finds the optimal order of the individuals for the social-services intervention. We normalize the input data using a robust nonlinear scaling that fits a logistic function to the empirical cumulative distribution of the data. We use a subsampling method to disregard from the training data 85% of the individuals whose level of involvement in crime is insignificant, so that the modeling effort focuses on those individuals who have at least a reasonably high level of risk. We also use a logarithmic transformation to scale the labels (correct answers for the predictions) so they have a comparable range of values and to increase the performance of the model.

Evaluation Method

We evaluate the performance of various candidate machine learning models by using nested, stratified, 5-fold cross-validation over the JacketID numbers to ensure unbiased separation of the training and testing data. Cross validation is a way of repeatedly dividing historical data into data sets used to *train* the model, and independent sets of data used to *test* the model fairly. This provides a principled way to evaluate model performance, and also a basis for optimization. We used an unpaired t-test to statistically compare the performance of different models, with a p-value less than 0.05 considered statistically significant.

Candidate Models

In small data sets, such as the one used in this project, linear models almost always outperform nonlinear models. In addition, linear models have the benefit of being easily interpreted by users, such as social workers and police, because a linear model is simply a weighted sum of the features ($a \times feature1 + b \times feature2 + \cdots$). Inspired by this project, we have applied this interpretability approach to predicting a patient's risk of cardiac death for the medical field, gaining positive attention in two recent papers.

Deep learning models, such as convolutional neural networks (CNN) and long short-term memory (LSTM) are helpful for image- and signal-processing problems, as well as time-series analysis, but are not relevant to applications such as the present one, so they were not considered.

Thus, we focused our modeling efforts on the following linear candidate models: least-angle shrinkage and selection operator (LASSO), support vector machine (SVM), and ridge regression. We also considered polynomial regression which, for low orders, can yield mildly nonlinear solutions. Among these candidates, the winning model was the LASSO.

Although the LASSO is linear, it is a sophisticated state-of-the-art approach. It uses an optimization strategy that produces *shrinkage*, i.e., it automatically shrinks the number of candidate features to identify a small subset of features that are most predictive. Mathematically speaking, the LASSO optimizes a weighted sum of an L2 prediction error with an L1 regularization term. The regularization parameter is determined automatically as part of the model optimization.

Most-Predictive Features

The LASSO automatically shrinks the original set of 23 (see Table 3) to just the eight variables shown in Figure 4, where each bar shows the weighting factor for each feature. Thus, the most predictive feature for estimating a person's future crime risk is the recent trend in their pattern of victimization and arrests (the slope of a line fit to their recent frequency of incidents). The other features are counts of various crime involvements, as well as the 1st- and 2nd-degree connections described earlier.

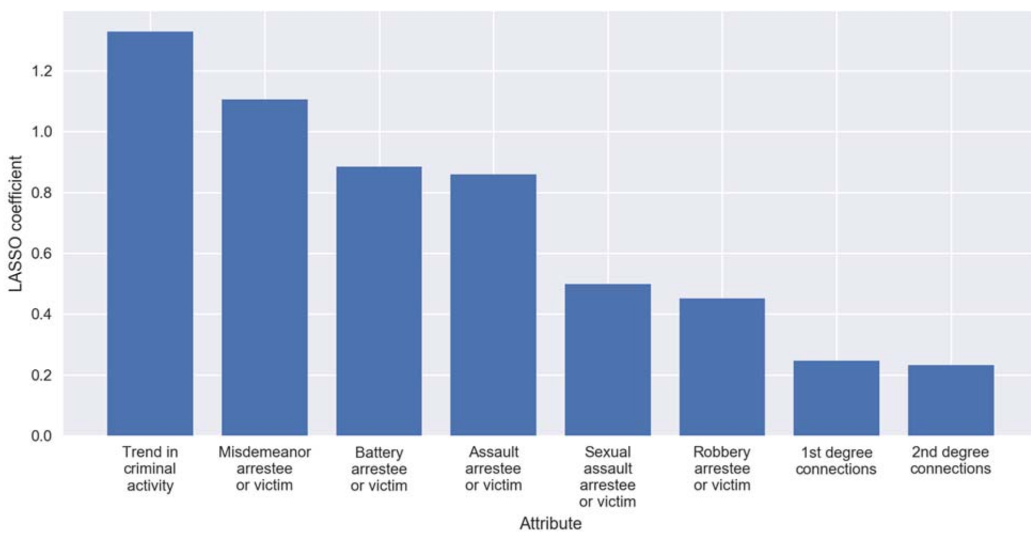

Figure 4. Variables found most predictive by LASSO in model Version 1.0.

Results

In Figure 5, the risk is represented by the red curve, showing that it drops rapidly as LASSO's scores decrease in magnitude. While the majority of the arrestee population have moderate risk (below 10% risk), those ranked highly by the algorithm are five times more likely to be involved in serious crimes (with probabilities around 50%); this shows that the model can differentiate people at high risk from the rest of the population.

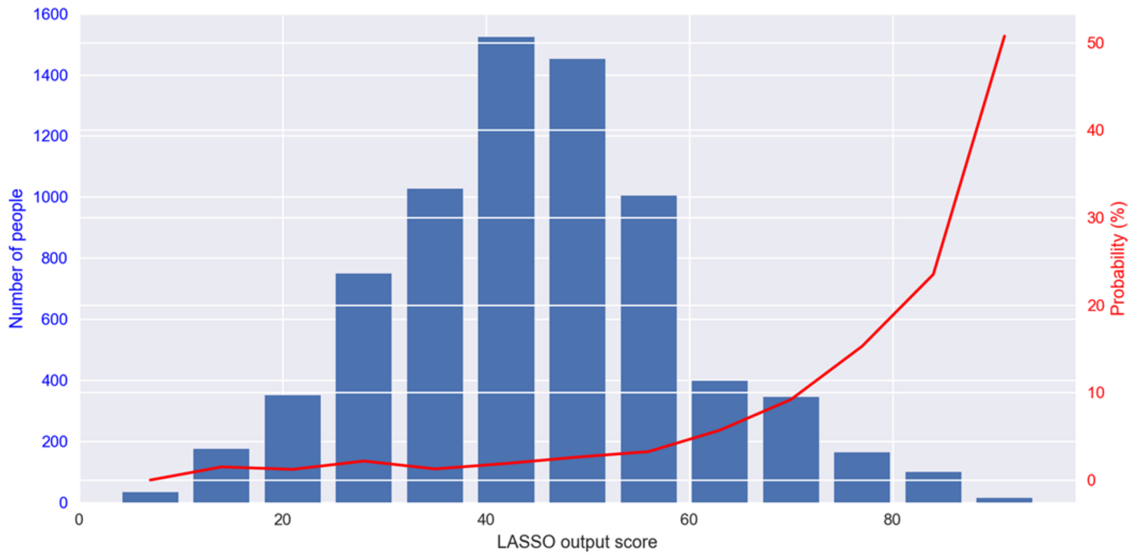

Figure 5. Performance of model Version 1.0. Blue bars are a histogram of risk scores produced by the model. The red curve shows the actual crime risk associated with a particular value of the risk score. This shows that about 20 people (the farthest right bar) have a greater than 50% chance of being involved in a serious crime in the next year (highest value on the red curve). The graph also shows that the vast majority of individuals have a crime risk below 5%. Thus, the model has successfully identified a very small number of people who may benefit most from social services intervention.

Figure 6 shows a further interpretation of the results of the model. This graph shows the total contribution to Elgin's crime severity for the *N* persons with highest projected risk. This shows that 30% of Elgin's overall crime issues are concentrated in a group of approximately 500 individuals who are identified by the model as having the greatest risk. This group of people represents only 6% of all arrestees in Elgin, and only 0.5% of Elgin's total population.

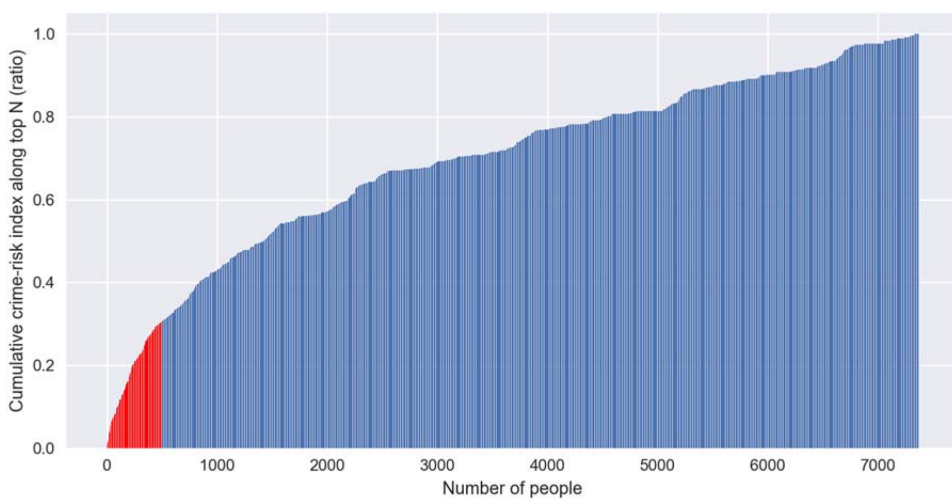

Figure 6. Normalized cumulative crime-risk index. The 500 individuals with the highest risk (in color red) capture 30% of the total crime severity.

Monitoring, Evaluating, and Implementing Risk Models and the Social Services Interventions

Appendix A is the official Standard Operating Procedure (SOP) for our social services intervention program, which is titled "Social Services Crime Prevention Program." This SOP document was developed through extensive collaboration between us and the EPD and was approved by the City of Elgin. The key aspects of the SOP are as follows:

- *The program will be based on social services only, and no law enforcement action will be taken against any individuals identified as at-risk by our model.* To be sure this is clear to the participants, they will be invited to participate by either the head social worker (Ms. Stingley) or the community liaison (Pastor Whitt), and not by a police officer. Meetings with participants will take place either at their homes by appointment, or in the social-services area of police headquarters.
- *The intervention program will be overseen by the head social worker (Ms. Stingley).* Ms. Stingley will follow up on each individual's case to ensure that good progress is being made, involving regular follow-up calls or visits with each participant.
- *Potential participants will be selected based on recommendations of the risk model, with final decision made by a committee including diverse members of EPD.* For example, the EPD may know that an individual is deceased, incarcerated, or no longer a resident of Elgin. The EPD may also know that the person has been consistently unwilling to participate in past social services programs.

Once the SOP was in place, we helped the EPD select a committee to review the information about the people identified by our model to be at highest risk. We attended the committee's first meeting, which followed the procedures that we helped the EPD develop for fair administration of the program.

That meeting revealed ways to improve our model and we have implemented them. Because of confidentiality concerns, the EPD provided us with a database that did not include names or personal information, only numeric identifiers (JacketIDs). Our assumption was that each JacketID was associated with one and only one identifiable individual. However, it turned out that some JacketID numbers were assigned to unknown suspects. We eliminated these unknown individuals from consideration.

EPD identified instances in their database in which a person was only a suspect of a crime and not a victim or arrestee. We then removed all such records from our analysis, and we have modified and evaluated the algorithm based on these changes. This step proved to be helpful to our model's success.

After the algorithm was revised, a subsequent meeting allowed the committee to assess the revised algorithm. The committee was surprised to learn that some of the people identified as being at highest risk were not the individuals that individual officers thought of as hardened criminals, but were instead homeless people. This has led to consideration of the types of social services that might be needed to best help that group.

Avoiding Racial Bias

During Year 2 of our project, a white police officer in Elgin was involved in a shooting of a female African-American suspect. Because of public attention focused on that shooting, we were asked by the police chief to provide information about the steps we have taken to assure that our model did not exhibit any racial bias. We were able to demonstrate scientifically and practically how we sought to minimize racial bias.

This issue is of critical importance since research on algorithms that predict future criminal activity (for purposes of deciding on issues such as bail or sentencing) have found that some algorithms make racially biased predictions. We decided at the start of our project not to use race or other variables overtly linked to race, such as zip code. However, bias may also exist in crime data due to the disproportionate arrests of people of color by police. Biases of this kind can be counteracted in various ways. Fortunately, the risk of bias is mitigated by the fact that machine learning algorithms seek quantitatively to maximize accuracy of the results, and bias is simply one form of inaccuracy that the algorithm works to remove. This is in contrast to human decision-makers, whose biases can lead to systematic judgment errors that go uncorrected. Finally, the risk of inadvertent racial bias can be further reduced by using historical data (as we did) to test the algorithm's output for the presence of bias and quantitatively seeking to minimize it.

We define model bias as follows. Suppose that individuals in a population belong to either racial/ethnic group A or B. A risk model is biased against group B if, among persons in groups A and B who are *predicted* to be at high and equal risk, the *actual* risk levels for group B are systematically lower than for group A. In other words, if we were to declare persons belonging to group B to be at high risk when actually they are not, then the risk model would be biased against group B.

To test our model, we looked at actual risks for individuals having the top 500 risk scores for three different time periods and computed the group means and standard deviations of the actual risks. We divided these individuals into three groups: white, Hispanic, and African American. (We omitted Asian/Pacific Islanders and American Indian/Alaskan Natives because the datasets contained too few individuals from these groups to permit analysis). For the persons with the top 500 risk scores, we found the mean actual risk was 23% for whites and Hispanics and 25% for African-Americans. This would imply a slight bias *in favor of* African-Americans, but the result was not statistically significant under any standard definition of significance (the p-values for all pairwise comparisons ranged from 0.40 to 0.74).

Based on our explanation of the design of our study, the Elgin Police Department was reassured that we were not perpetuating racial bias in our predictions.

Meeting the Goals of the Nayar Prize Program

The Nayar Prize program was established to encourage collaborations within the Illinois Tech community "to develop breakthrough, innovative projects that will, within three years, produce meaningful results with a societal impact." We chose to work in the area of crime prevention because of the profound effect crime has on the individuals involved, on their neighborhoods, on the business community, on government, and on the economy. Our interdisciplinary approach to the problem has provided an important contribution to the field by creating a way to predict who is at the highest risk of being involved in serious crime and to offer social services to help avert that risk.

We have also contributed to the national and international dialogue on the subject and brought attention to our project and enhanced the reputation of Illinois Tech. One of our group's crime risk models was the subject of the January 7, 2018 episode of "60 Minutes." The episode covered an actual instance where a 31-year-old African American gang member and former drug dealer was identified by our predictive policing algorithm as high-risk and was subsequently enrolled into a social services program which helped him get a state ID and find employment. He praised our approach, noting the appreciation he felt for being able to work his first job, a part-time position at an animal kennel: "Man, it was like heaven, you know? Even though I was a drug dealer, you know, like, I always, kind of, had money, but it feels different when you work for it. I want to keep working. I don't ever want to go back to the streets."

We have been involved in two meetings with the National Courts and Sciences Initiative (NCSI), a non-profit organization that trains judges to evaluate scientific and technological issues in court. Because of these meetings, the NCSI has approached us to serve as a resource for judges nationwide and to host a national meeting of judges in 2019 at Chicago-Kent College of Law, where we will train judges to assess algorithms used in the criminal justice system. We will also present the results of our Nayar Prize study.

In the Spring 2017 and Spring 2018 semesters, Miles Wernick participated as a guest speaker on the subject of this project in an undergraduate Illinois Tech class. In addition, Lori Andrews was a plenary speaker at an international gathering of lawyers and social scientists in Toronto in June 2018, where she additionally participated in sessions on predictive policing. Later in the same month, Lori lectured prosecutors and police officials from Barbados, the Czech Republic, Hong Kong, Mexico, and Nigeria about the Nayar Prize crime prevention intervention project. She also described the project at a session of the American Bar Association Annual Meeting (which brings together thousands of lawyers from around the world). We have also given talks about our approach at other universities—including University of California Irvine—and to companies, including the financial technology company Remitly in Seattle.

Plans for Year 3

Study Impact of Data-Driven Intervention

In Year 3, we will continue to enhance our model and critically examine the success of our data-driven intervention techniques. We expect to spend considerable time with the social workers learning about which interventions have been most helpful to which people. Social workers will be provided with a standardized reporting form that tracks such information as the dates of interactions with at-risk individuals, the nature of the interaction, what information was gathered, what action was taken, and what follow-up was planned. EPD will add to these records any crime incidents in which the individual was a victim or arrestee, including the date and nature of the incident, and the individual's role. This ongoing record will provide data to observe the outcomes of the interventions.

To rigorously measure the effect of the social workers' intervention would involve tracking the program for a number of years; however, the data gathered during this pilot study will provide evidence suggesting what works and does not work, and will give clues as to whether the model's recommendations are sound in the practical setting of social services. It will also provide sufficient data to permit a rough power analysis to be performed that would be a necessary pre-requisite for designing a future larger-scale trial. We will also be able to obtain feedback from social workers regarding the usefulness of our model's recommendations versus those obtained though chance encounters, which is the typical way in which social workers are alerted to at-risk people.

Analyze Relationship between Homelessness, Mental Illness, and Crime

According to our risk model, homeless people have an especially high risk of being involved in a serious crime in Elgin. In Year 3, we will test the accuracy of this prediction against empirical studies of the correlation between homelessness and crime. Since mental illness is widely perceived as a primary risk factor for homelessness, we will also investigate how a co-occurrence of mental illness and homelessness may be connected to serious crime. Various studies have found that, in general, homeless people are more likely than domiciled people to commit non-violent crimes, but are less likely to commit violent crimes. Over the next year, we will compare these findings to those of our risk model as we continue to improve the model's predictive power.

Explore the Possibility of Expanding to Other Communities

The success of our model in Elgin has encouraged us to explore expanding the crime prevention program to other communities, such as Rockford, IL. This could provide the opportunity to monitor the results of our program in two communities and evaluate how the effectiveness of our intervention techniques changes from one community to another.

A Special Note

David Haro Alonso received his PhD from IIT in May 2018 based in part on his research for this project. His dissertation was titled "Individual-Based Risk Models for Crime Prevention and Medical Prognosis." For his dissertation, David studied risk-assessment models used to make predictions and to categorize individuals according to their likelihood of an adverse outcome. He developed such models for two real-world problems in crime prevention (in collaboration with the Chicago and the Elgin Police Departments) and one in preventive medicine (with the Cedars-Sinai Medical Center from Los Angeles, CA).

Attachments

Appendix A: Standard Operating Procedure (SOP)

	ELGIN POLICE DEPARTMENT 151 Douglas Avenue Elgin, Illinois 60120	
Effective Date:	**STANDARD OPERATING PROCEDURE**	**Revised Date:**
Chief of Police:	Social Service Crime Prevention Program, 42.9	
Cross Reference:	**Policy Sections:** 42.9.1 Program Administration 42.9.2 Program Selection 42.9.3 Department Review & Eligibility for Services 42.9.4 Intervention 42.9.5 Program Evaluation	

PURPOSE

The purpose of this policy is to establish guidelines for the administration of the department's Social Service Crime Prevention Program (SSCPP).

POLICY

It is the policy of the Elgin Police Department to endeavor to create a safer community. The SSCPP is a collaboration with the Illinois Institute of Technology (IIT) to identify individuals with a high propensity for being a party to a serious crime, in the capacity of a victim or offender. The department, in conjunction with participating social service agencies, aims to provide the appropriate services to program participants to prevent his/her involvement with a future crime.

DEFINITIONS

Risk model: A mathematical procedure based on statistics that estimates the chances of something happening in the future based on risk factors. For example, smoking is a risk factor for cancer, because it increases the chances. Similarly, a consistent pattern of involvement in crime incidents is a risk factor for a person's continued involvement.

PROCEDURES

42.9.1 PROGRAM ADMINISTRATION

A. The program will be administered by a supervisor designated by the chief.

B. The head department social worker will oversee the intervention portion of the program.

C. The head social worker will report on the status of the intervention on each person in the program to the program administrator at least bi-weekly.

D. The program administrator will track all persons in the program and determine when additional people shall be added. A person will generally be added when a person currently in the program has completed services, refuses initial services or ends participation.

E. A department information technology specialist will work with the program administrator and IIT on any technical needs for the program.

42.9.2 PROGRAM SELECTION

A. The department will provide data to IIT personnel for the exclusive purpose of creating a risk model that will help guide the selection of participants who might from the SSP. The data that will be provided to IIT are confidential and IIT personnel shall not release or utilize the data for any other purpose.

B. The data may include, but are not limited to arrests/bookings, incidents, criminal reports, and parole status. The data provided to IIT will not include a person's name, race, ethnicity, physical characteristics or place of residence to avoid potential of bias in the results (except to the extent that, racial data may be used apart from the risk analysis to double-check that the algorithm has no racial bias). All persons provided in the data set will be identified by their unique jacket number, as listed in the department's records management system.

C. IIT will develop a risk model that will identify individuals having a high risk of being involved in a serious crime so that the social service interventions can be focused on individuals who most need it.

D. The results of IIT's risk analysis will be forwarded to the program administrator.

E. The analysis will contain the following information:
   1. An estimated risk level for each individual with a recent arrest history.
   2. An explanation of how each individual's risk factors led to the estimated risk.
   3. The relative severity (compared to other individuals) of each of the individual's risk factors, expressed as a percentile

42.9.3 DEPARTMENT REVIEW AND ELIGIBILITY FOR SERVICES

A. A review of the data set will be conducted by a committee composed of diverse members of the department with varying job responsibilities. The review process will be coordinated by the program administrator. The purpose of this review is to ensure that the individuals identified in the analysis are appropriate candidates for the program. The review shall consider, but not be limited to, the following information:

1. Current custody status with the Illinois Department of Corrections (IDOC).
2. Residency: The individual resides in the city limits.
3. Pending or recent police contact, to include types of contact.
4. Likelihood of participating in the program.

B. Records shall be maintained for those individuals not selected; these records shall provide a brief explanation as to why the person was not selected to participate.

C. Individuals under the age of 18 are not eligible for this program. Furthermore, data pertaining to juveniles is excluded from the data provided to IIT.

42.9.4 INTERVENTION

A. Initial contact with selected individuals will generally be conducted by a sworn officer and a social worker or department community outreach specialist. When possible, the initial contact will be made at the individual's home, by appointment set up by the social worker. If the individual wishes to meet at the station, the meeting should be conducted in the social services area or similar type room.

B. During the initial contact, the individual should be given information about the program and the desire for the department to facilitate services to reduce the likelihood of involvement in criminal activity. An appointment should then be scheduled for the subject to meet with a department social worker for an initial evaluation for services.

C. The head social worker shall assign a social worker to work with the subject. The program administrator will also assign an officer who will serve as a liaison to the individual.

D. During the initial evaluation, the assigned social worker and the individual shall determine which of the available services the individual might need to help prevent future involvement in a serious crime, either as a victim or as a perpetrator. These services may include education, job training, mental health services, and substance abuse treatment. Referrals shall then be made to the appropriate partner agencies.

E. Following the referrals, the social worker and liaison officer shall make regular follow-up calls or visits with the individual to track progress and assist with services.

F. The liaison officer shall also make a designation in the department records management system to track any police contacts with the subject.

42.9.5 PROGRAM EVALUATION

A. Quarterly, the program administrator will provide a report showing the progress of each individual selected in the program to the deputy chief. The report will include progress on social service utilization and any criminal activity involvement by the individual.

B. The completed report will be provided to IIT, who will conduct an evaluation of the program to determine its effectiveness and to provide further insight on any modifications suggested improvement.

Appendix B: Courts' Determination of the Permissibility of a Factor a Police Officer Used to Predict That an Individual Encountered in Public Would Commit or Had Committed a Crime

	Permissible	Not Permissible	Total	Percentage Permissible
Driving				
Not Suspicious On Its Face				
Broken brake light	1	0	1	100%
Does not immediately pull over	1	0	1	100%
Does not use turn signal	1	0	1	100%
Driving fast (not speeding)	1	0	1	100%
Driving on shoulder	1	0	1	100%
Driving with multiple occupants visible	2	0	2	100%
Driving without seatbelt	2	0	2	100%
Marijuana scent	1	0	1	100%
Out of state license plate	1	0	1	100%
Sports car	0	1	1	0%
Possibly Suspicious				
Matches car description of wanted/suspected criminal	3	0	3	100%
Problem with license/registration	2	0	2	100%
Related to Something Clearly Illegal				
Speeds	2	0	2	100%
Body Language				
Not Suspicious On Its Face				
Argues/Shouts	2	0	2	100%
Bulge (possibly gun) visible	2	0	2	100%
Drinks in public	1	0	1	100%
Fidgets hands (possibly intoxicated)	1	0	1	100%
Grabs waistline	1	0	1	100%
In the middle of the street	2	0	2	100%
Keeps hands in pockets	1	0	1	100%
Looks around (possibly for escape option)	1	0	1	100%
Sways (possibly intoxicated)	1	0	1	100%
Wears bandana over face	2	0	2	100%
Nervous	11	1	12	92%
Looks/Checks something	1	1	2	50%
Gives the middle finger to police	0	1	1	0%

Paces	0	1	1	0%
Rubs head	0	1	1	0%
Possibly Suspicious				
Holds/Reaches for object (possibly illegal substance)	1	0	1	100%
Holds/Reaches for object (possibly weapon)	9	0	9	100%
Holds/Reaches for object (unsure what it could be)	4	0	4	100%
Race				
Not Suspicious On Its Face				
Inability to speak English	2	0	2	100%
Middle Eastern-looking	1	0	1	100%
Speaking Spanish	1	0	1	100%
Apparent Hispanic/Latino ancestry	1	1	2	50%
Black	0	7	7	0%
Possibly Suspicious				
Matches race of known/suspected criminal	16	0	16	100%
Public Appearance				
Not Suspicious On Its Face				
Cell phone battery dead	1	0	1	100%
Has no key to own apartment	1	0	1	100%
In a high-crime area	27	0	27	100%
In a parking lot	14	0	14	100%
Near a closed building	1	0	1	100%
Near an ATM	1	0	1	100%
Out in the cold	2	0	2	100%
Scares others	1	0	1	100%
Seems out of place/does not mingle with others	1	0	1	100%
Sits in car	9	0	9	100%
Wears baggy clothing	1	0	1	100%
Out late at night	11	1	12	92%
Has a beeper	0	1	1	0%
In a predominately white neighborhood	0	1	1	0%
Reads newspaper	0	1	1	0%
Wears expensive jewelry	0	1	1	0%
Possibly Suspicious				

Has electronics in possession that match description of stolen goods	1	0	1	100%
Matches description (non-racial description)	4	0	4	100%
Near a recently reported crime	12	0	12	100%
Taken Actions				
Not Suspicious On Its Face				
Avoids contact with police	10	0	10	100%
Difficulty remembering information when questioned by authorities	2	0	2	100%
Complies with police	1	1	2	50%
Possibly Suspicious				
Does not comply with police	5	0	5	100%
Runs/Drives away quickly	8	0	8	100%
Tries to hide	5	0	5	100%
Unable to/Does not provide identification or proof of citizenship	2	0	2	100%
Uses hands to shuffle around other objects to hide something	1	0	1	100%
Related to Something Clearly Illegal				
Walks toward courthouse with gun	1	0	1	100%
With Others				
Not Suspicious On Its Face				
Inconsistent stories	2	0	2	100%
Others leave quickly	2	0	2	100%
Others try to hide	2	0	2	100%
Splits up from others quickly	1	0	1	100%
With suspicious people	5	2	7	71%
Possibly Suspicious				
With known prostitute	1	0	1	100%

References

[1] David Haro Alonso, Miles N. Wernick, Yongyi Yang, Guido Germano, Daniel S. Berman, Piotr Slomka, Prediction of cardiac death after adenosine myocardial perfusion SPECT based on machine learning, *Journal of Nuclear Cardiology*, 1-9, 2017.

[2] S. Shresthra and P.P. Sengupta, Machine learning for nuclear cardiology: The way forward, *Journal of Nuclear Cardiology*, 1-4, 2017.

[3] J. Walker, Roskam, Elgin police honor department's 'unsung hero', *Chicago Tribune*, March 9, 2016.

[4] J. Walker, Personal, institutional racism focus of Black Lives Matter discussion, *Chicago Tribune*, February 26, 2016.

[5] *Wisconsin v. Gallion*, 678 N.W.2d 197, 209 (Wis. 2004).

[6] *State v. Loomis*, 881 N.W.2d 749, 770 (S. Ct. Wis. 2016).

[7] *Hall v. Pennsylvania State Police,* 570 F.2d 86 (3d Cir. 1978).

[8] *Whren v. United States*, 517 U.S. 806, 810 (1996).

[9] *Whren v. United States*, 517 U.S. 806, 813 (1996).

[10] *Buck v. Davis,* 137 S.Ct. 759 (2017).

[11] *Gardner v. Florida,* 430 U.S. 349, 352 (1977).

[12] *State v. Pfenning*, No. 57-4-96 Grand Isle County, Vermont (Apr. 6, 2000) (unpublished decision).

[13] 2016 Law enforcement use of social media survey, *International Association of Chiefs of Police* (February 2017).

[14] K. Leetaru, Geofeedia is just the tip of the iceberg: the era of social surveillance. *Forbes* (October 12, 2016).

[15] M. Cagle, Facebook, Instagram, and Twitter provided data access for a surveillance product marketed to target activists of color. *American Civil Liberties Union Blog* (October 11, 2016).

[16] M. L. Williams, A. Edwards, W. Housley, P. Burnap, O. Rana, N. Avis, J. Morgan, L. Sloan, Policing cyber-neighbourhoods: tension monitoring and social media networks. *Policing and Society* **23**, 1-21 (2013)

[17] *In re Warrant to Search a Target Computer at Premises Unknown*, 958 F.Supp.2d 753 (S.D. Tex. 2013).

[18] T.-Y. Liu, *Learning to Rank for Information Retrieval*. Springer Science & Business Media, 2011.

[19] K. Järvelin and J. Kekäläinen, IR evaluation methods for retrieving highly relevant documents, in *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 41-48, 2000.

[20] Y. Wang, L. Wang, Y. Li, D. He, W. Chen, and T.-Y. Liu, A theoretical analysis of NDCG ranking measures, in *Proceedings of the 26th Annual Conference on Learning Theory (COLT 2013)*, 2013.