ILLINOIS INSTITUTE OF TECHNOLOGY
impact

Nayar Prize II Phase I Quarterly Progress Report (Q3)
July 2017

Project: A Data-Driven Crime Prevention Program
Team: Lori Andrews, Miles Wernick, and Yongyi Yang
Students: David Haro Alonso, Pauline Panay, and Raven Zeer

Progress Summary

The goal of this project is to design, implement, and deploy a flexible, new model for crime prevention in Elgin, Illinois. The model can be translated to a wide array of communities in the U.S. and beyond, thereby achieving far-reaching societal impact. Our approach uses a computerized prediction model to help identify individuals at risk for significant involvement in crime so that social-services assistance can be offered to help them improve their lives. In the third quarter of the project (Q3), we refined and evaluated the prediction model, and used legal and policy analysis to ensure that the model is sensitive to ethical, legal, and social concerns. Already we have been asked by other scholars and service providers for access to our model and we are developing additional partnerships to expand its reach and impact.

Predicative Model Development

In Q3, we developed model version v0.3, which predicts the total "severity" of crimes in which an individual will be involved in the upcoming year, where we now define "severity" of each crime based on Illinois' prescribed maximum sentence (in years). This resulted in an algorithm that provides excellent results. The graph at right shows the histogram of predicted risk scores for recently arrested individuals (blue bars) along with their actual risk of being involved in a serious crime (red curve). This shows that the model successfully identifies a very small number of individuals who have about a 46% chance of involvement in serious crime in the upcoming year, which is very high, even compared with other recent arrestees. Experts at the Elgin Police Department (EPD) have begun providing feedback on the model's recommendations, and initial feedback has been immensely positive. If further review identifies key areas for improvement, we will conduct a further round of

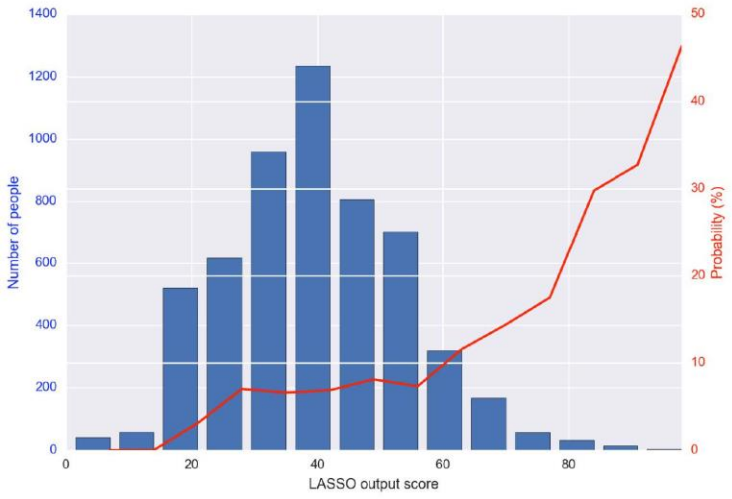

model refinement; if not, then v0.3 will be designated as v1.0, which will be the version deployed in the initial interventions.

Assuring Legal, Ethical, and Social Appropriateness

Because of the high societal impact of our work, it is critical to analyze relevant laws, policies, and ethical guidelines to ensure that our new program meets the highest standards. Such analyses will have an impact beyond the Elgin project by educating the public, law enforcement, and social services agencies as to the appropriate uses of our approach, while identifying potential pitfalls to be avoided. To that end, we analyzed 100 relevant court decisions about police decision making, analyzed all studies of predictive policing attempts (most of which attempt to predict where crime will occur, rather than who will be involved), and assessed media coverage of predictive policing. Based on this research and on sound algorithm design principles, we have identified the following relevant policies that should be taken into consideration in the development of predictive policing models. These will provide a valuable resource to the field.

First, the algorithm must be publicly available. If decisions are made about an individual based on a secret algorithm, this can infringe upon the person's constitutional rights. For example, 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. Other court decisions have said that using secret algorithms "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…claims." Of course, we will be publishing our algorithm.

Second, the algorithm must not use variables that could lead to racially-discriminatory results. A computerized risk model can 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 2017, the U.S. Supreme Court underscored the impermissibility of using race to predict dangerousness. In Buck v. Davis, the defendant's expert said that race is a factor "know[n] to predict future dangerousness." The Court wrote, "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." In our algorithm, we have eliminated from consideration any variable that might be indicative of race, and we conduct further testing to ensure that racial bias has not inadvertently entered into the results.

Third, the algorithm must not use variables in ways that infringe rights of free speech, association, and privacy. It is appropriate to use public information such as arrest records, because these indicate there has already been a reasonable societal suspicion about the individual's connection to criminal behavior. On the other hand, the use of a variable based on social media posts or surreptitious surveillance would be scientifically unsound and legally improper. Yet such services are currently

available. For example, a service called SnapTrends monitors tweets and social media posts for signs of dangerousness. Proponents of the service suggest that, for example, the police should investigate someone who has posted the sentence “I’ve had enough.” This violates a person’s right of freedom of expression and freedom of association, and that of friends and relatives who communicate with him on social media.

Finally, the algorithms should not be used by the criminal justice system in punitive ways such as for sentencing. Algorithms designed to predict future violence should be used primarily to allow for interventions to prevent future crime by helping the individual. They should not be used as a justification to keep people under surveillance, determine their sentences, or set conditions on parole. 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.’”

Our Strategy

The overall goal of our project is to demonstrate that technology can be used together with social assistance programs to help at-risk individuals and make communities safer. By adhering to the legal, ethical, and social principles outlined above, we hope that our project will set a positive example for others to follow. In Q4 we will also begin preparing publications to advocate these principles and thereby contribute to society’s discussion of strategies for crime prevention and the appropriate role of technology.