TEMPLE UNIVERSITY
Center for Security and Crime Science

THE PHILADELPHIA PREDICTIVE POLICING EXPERIMENT

EFFECTIVENESS OF THE PREDICTION MODELS

Jerry H. Ratcliffe, Ralph B. Taylor and Amber P. Askey

Overview

The Philadelphia Predictive Policing Experiment represents a two-year collaboration between Temple University's Center for Security and Crime Science, housed in the Department of Criminal Justice at Temple, and the Philadelphia Police Department. This National Institute of Justice-funded research project was the first place-based, randomized experiment to study the impact of *different* patrol strategies on violent and property crime in predicted criminal activity areas. The experiment aimed to learn whether different but operationally-realistic police responses to crime forecasts estimated by a predictive policing software program can reduce crime.

Software performance

The predictive policing software used was the HunchLab program designed by Azavea. HunchLab is a web-based predictive policing system that accesses real-time Philadelphia Police data to produce crime forecasts for the city. It incorporates statistical modeling that considers seasonality, risk terrain modeling, near repeats, and collective efficacy. Officers at police district buildings in the marked and unmarked conditions could log in and print out maps for forthcoming 8-hour shifts. Azavea adapted the software at the request of the Philadelphia Police Department and researchers from Temple University to generate three predicted 500 feet square grids per district per shift. They also included a slight randomization component to reduce the possibility that the same grid cells were predicted every day. This allowed the project to create the necessary experimental conditions by limiting the output of the software to generate three predicted grids (500 feet by 500 feet) in every district of the city each day. **It is important to note therefore that the experiment artificially reduced the efficiency of the software**, because it forced the software to choose grids in low crime districts, and limited the number of grids it could assign in high crime districts.

During the experiment, researchers accompanied officers on over 100 dedicated car assignment patrols (either marked or unmarked cars). It was clear from these observations that it was impossible and unrealistic for assigned officers to exclusively patrol the 500 feet square predicted grid areas. In many cases the grids identified one-way streets, necessitating driving through the surrounding block or two to return to the grid assignment. We therefore examine the efficacy of the software considering not just the predicted grids

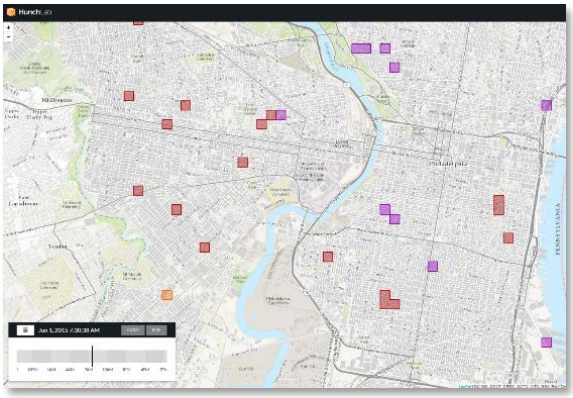

alone, but also and more realistically the surrounding grids as well. We do this for the grids nearest to each predicted grid (in technical terms, *first order queen contiguity*), and for grids within two cells of each predicted grid (*second order queen contiguity*). The table below shows the results of the software prediction for control districts, as the control districts were not affected by focused police activity resulting from the experiment. The control districts are a good indication of the efficacy of the software program.

The unit of analysis is the district day – that is, an eight-hour shift in one district in one day. For each experimental condition (such as the control sites shown here) there were five districts across a 90 day experimental phase. This equals 450 district days (5 districts x 90 days = 450) and 460 days for the violent crime phase (the violent crime phase ran for two additional days, so 5 districts x 92 days = 460).

In the table you can see that in the control sites for property crime, 274 district days (60.9%) had at least one crime in the entire district during the shift. For property crime in these 274 district days, the software predicted nearly 14 percent of the crime available to predict in these districts, and did so by identifying just 6.8 percent of the district area. When extended up to two grid cells away from the predicted grids, it predicted 30.3 percent of the property crime while highlighting only 15.3 percent of the district. Please remember that these predictions were not optimized – the software was deliberately hobbled to enable the experimental design. The results are therefore conservative. Also note that the results differed by district and crime type, as shown below. **In general, however, it appears that the software was able to predict twice as much crime as we would expect if crime were spread uniformly across the districts, even when artificially constrained by our experiment to be less effective than designed.**

Phase / district	**District days**	**Crime days in district**	**Total crimes**	**Within one grid (% total)**	**% area of district**	**Within two grids (% total)**	**% area of district**
Property	**450**	**274 (60.1%)**	**489**	**68 (13.9%)**	**6.8%**	**148 (30.3%)**	**15.3%**
24	90	61 (67.7%)	127	11 (8.6%)	5.1%	31 (24.4%)	11.4%
16	90	46 (51.1%)	64	12 (18.7%)	6.1%	22 (34.4%)	12.9%
35	90	74 (82.2%)	157	22 (14%)	5.7%	34 (21.6%)	14.2%
9	90	59 (65.5%)	92	15 (16.3%)	12.9%	38 (41.3%)	28.6%
1	90	34 (37.7%)	49	8 (16.3%)	2.6%	23 (46.9%)	4.9%
Violent	**460**	**137 (29.8%)**	**176**	**24 (13.6%)**	**5.1%**	**43 (24.4%)**	**12.1%**
25	92	49 (53.2%)	68	7 (10.3%)	5.8%	18 (26.5%)	13.9%
35	92	32 (24.4%)	40	4 (10%)	5.9%	6 (15%)	14.4%
26	92	19 (20.6%)	20	4 (20%)	7.4%	10 (50%)	16.2%
14	92	27 (29.3%)	37	8 (21.6%)	2.6%	8 (21.6%)	6.5%
8	92	10 (10.8%)	11	1 (9.1%)	1.4%	1 (9.1%)	3.6%

Further information

For additional and current information, please visit the project website at **bit.ly/CSCS_3PE**

Suggested citation: Ratcliffe, JH, Taylor, RB and Askey, AP (2017) *The Philadelphia Predictive Policing Experiment: Effectiveness of the Prediction Models*. December 2017. Online at bit.ly/CSCS_3PE. 2 pages.