

Rewiring police officer training networks to reduce forecasted use of force

Ritika Pandey

Department of Computer and Information Science,
Indiana University - Purdue University
Indianapolis, IN
ripande@iupui.edu

James Hill

Department of Computer and Information Science,
Indiana University - Purdue University
Indianapolis, IN
hilljh@iupui.edu

Jeremy Carter

O'Neill School of Public and Environmental Affairs,
Indiana University - Purdue University
Indianapolis, IN
carterjg@iupui.edu

George Mohler

Department of Computer Science,
Boston College
Chestnut Hill, MA
mohlerg@bc.edu

ABSTRACT

Research has shown that police officer involved shootings, misconduct and excessive use of force complaints exhibit network effects, where officers are at greater risk of being involved in these incidents when they socialize with officers who have a history of use of force and misconduct. In this work, we first construct a network survival model for the time-to-event of use of force incidents involving new police trainees. The model includes network effects of the diffusion of risk from field training officer (FTO) to trainee. We then introduce a network rewiring algorithm to maximize the expected time to use of force events upon completion of field training. We study several versions of the algorithm, including constraints that encourage demographic diversity of FTOs. Using data from Indianapolis, we show that rewiring the network can increase the expected time (in days) of a recruit's first use of force incident by 8%. We then discuss the potential benefits and challenges associated with implementing such an algorithm in practice.

CCS CONCEPTS

• Applied computing → Law, social, and behavioral sciences.

KEYWORDS

Survival analysis, network optimization, annealing, police use of force

ACM Reference Format:

Ritika Pandey, Jeremy Carter, James Hill, and George Mohler. 2023. Rewiring police officer training networks to reduce forecasted use of force. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3580305.3599899>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0103-0/23/08...\$15.00

<https://doi.org/10.1145/3580305.3599899>

1 INTRODUCTION

In the wake of George Floyd's murder on May 25, 2020 at the hands of Minneapolis police officer Derek Chauvin, large protests across the United States called for police to be reformed, defunded, or abolished. The officer ultimately convicted of second degree murder had 18 prior conduct complaints, 2 of which resulted in reprimands [22]. Furthermore, 3 other officers (two of whom were trainees) were present and did not intervene when George Floyd said "I can't breathe" [20]. At the time of writing the present article in January 2023, police excessive use of force remains a significant problem in the U.S.; on January 7, 2023, Tyre Nicols was killed by five police officers in Memphis, TN during a traffic stop, where video shows that he did not fight back [9]. While sweeping changes are needed to address systemic racial biases in the criminal justice system, from traffic stops to incarceration, an immediate reform that may be achievable in the short term is improving officer risk assessments and interventions.

Recent research has shown that officers with negative marks on their record (complaints, firearm discharges, etc.) are 3 times more likely to shoot in the line of duty [23]. Other research has shown that officer shootings, misconduct and excessive use of force exhibit network effects, where officers are at greater risk of being involved in these incidents when they socialize with officers who have a history of misconduct and complaints [21, 27, 32]. Given that use of force and misconduct behavior appear to be transmissible across police networks, this leads to the question: can police networks be altered to reduce use of force and misconduct events?

In the present article we attempt to take a small step towards answering this question. We analyze field training data from Indianapolis, Indiana where new recruits have been paired with field training officers (FTO). We first build a survival model for the time to the first use of force incident of a new police officer after field training, where the survival model includes covariates on both individual level characteristics of the officer and FTO characteristics to capture network effects (such as use of force history of the FTO). We then apply a simulated annealing algorithm to the fitted survival model, to rewire the network in order to maximize the average time-to-event of use of force across the network. Comparing the results of the optimized network to the network used by IMPD from 2018

to 2021, we find that our rewiring algorithm increases the forecasted (average) time to use of force events by 8%.

The outline of the paper is as follows. In Section 2, we discuss related work on network intervention algorithms, survival modeling, and studies that investigate the impacts of field training on use of force. In Section 3, we discuss the field training process and data used in our study from the Indianapolis Metropolitan Police Department (IMPD). In Section 4, we construct a network survival model for the time to use of force events and then introduce a novel simulated annealing algorithm for optimizing the network. In Section 5, we discuss results from several experiments where our model was applied to IMPD data. In Section 6, we discuss the potential benefits and challenges associated with implementing police network rewiring algorithms in practice.

2 RELATED WORK

Network interventions may target individual or groups of nodes, encourage or discourage certain peer-to-peer interactions to change information/behavioral diffusion, or alter the network through node deletion or rewiring [26]. Node-level interventions can be used for a wide range of tasks, from influence maximization [2] to determining optimal immunization strategies [31]. Simple and complex contagions can also be controlled through edge removal or deletion [10, 11]. Networks can be modified to optimize key network statistics like group centrality [19] or reinforced to make them more resilient during disasters [30]. Rewiring the entire network is a less studied problem compared to node/edge removal algorithms, as one typically does not have control over the entire network. However, in some situations this is possible, for example network rewiring has been applied to what-to-watch-next recommendations [4]. More recently, network rewiring has been studied in the context of adversarial attacks on graph neural networks [15].

Survival analysis has been applied in a number of data mining contexts. In [12], a survival model was applied to forecast time-to-talent turnover events and career progression. In [14], survival analysis was applied to Reddit posts to predict transitions to and from drug addiction. Survival analysis has also been applied in the context of information diffusion on networks, for example to evaluate the rate of adoption of public health interventions [29].

Several past studies have examined the direct impacts of training on officer use of force or misconduct. In [13], the authors examine the effect of supervisor education and training on police use of force practices and, in [5], the relationship between police field training officers (FTOs) and their trainees' subsequent allegations of misconduct were investigated. To date no studies have considered the problem of altering the FTO-trainee network to reduce forecasted use of force.

The contribution of the present paper is two-fold. 1) To our knowledge there is no existing work that considers optimizing survival models on networks through rewiring. The algorithm we introduce here fills in this gap and applies to a variety of time-to-event problems on networks where the goal is to rewire the network to either increase or decrease the expected time-to-event. 2) Our model is applied to the important problem of reducing the risk of use of force by rewiring police social networks. We study a novel police training social network dataset, and propose an algorithm for rewiring such




Figure 1: Number of failed and right-censored officers in the different risk sets (during training, 6 months and 12 months after training). Here failure time is the time of occurrence of a use of force incident.

a network early in the career of officers, when such an intervention may be most effective.

3 POLICE FIELD TRAINING DATA

Field training is formalized, on-the-job instruction for new police recruits. After being hired, recruits receive training over a several month period at a training academy, where they are provided with classroom instruction on criminal justice, human behaviour and police skills. After graduating from the academy, recruits are paired with field training officers (FTOs) who serve as mentors, trainers and evaluators. Field training provides exposure to field problems and patrol situations, where new police officers can apply classroom principles to real-life situations [6, 16, 25].

The data in the present study consists of field observations and use of force incidents for Indianapolis Metropolitan Police Department (IMPD). IMPD is the law enforcement agency serving Indianapolis, Indiana and employs approximately 1700 officers. This research is based on three data sources: field training data collected from recruit classes during December 2018 to January 2021, officer use of force incidents, and roster data on all officers who were either FTOs or trainees during this time period.

We note that use of force can be sub-classified into several categories: 1) legal, if applied in accordance with the law, 2) excessive, when an officer uses more force than necessary or legal, and 3) police brutality, where no actual law enforcement function is occurring and the officer is being abusive. In the present study we focus on aggregated use of force incidents where we do not distinguish between these three types of events, in part because few incidents in the dataset are labeled as excessive.

The training period in Indianapolis is 5 months, followed by a 1 year probationary period. Table 1 summarises officer level attributes in the dataset. For predicting an officer's time to their first (or next) use of force, we separate our analysis into 3 different risk sets—**during training, 6 months after training and 12 months after training**. In the first training period, the forecast is made at day 0 and the observations are collected through the 5 month training period. In the second period, the forecast is made at the end of the 5 months of training and lasts for 6 months (11 months after the start of training). In the third period, the forecast of the time to the next

use of force incident is made at the end of the 5 months of training and lasts for 12 months (17 months after the start of training).

A key advantage of survival models over traditional regression models is their ability to effectively handle right-censored data. Right-censored data refers to observations in which the event of interest has not yet occurred, such as an officer who has not had an incident by the end of the observation period. Traditional regression models assume that all data points are fully observed and cannot account for the censoring that can occur due to a finite observation window. Figure 1 depicts the failure-censored group distribution of the 189 recruits in our study across the three different risk sets.

Table 1: Officer Attributes of IMPD Field Training from Dec 2018 to Jan 2021.

Attribute	FTO	Trainee
N	240	189
Gender		
Male	85%	79.9%
Female	15%	20.1%
Race and ethnicity		
White	86.7%	74.6%
Black	10.8%	16.4%
Hispanic and other	2.5%	9.0%

4 METHODOLOGY

In our network survival model, each officer in training is represented by a feature vector that includes trainee sex, race, assigned policing district (location), and average evaluation scores during training. IMPD jurisdiction is divided into north, east, northwest, southeast, and southwest districts. Due to high correlation, southwest and southeast were combined into a single grouped location feature. Field trainees are scored daily by FTOs on categories such as appearance, attitude, knowledge, and performance on a scale of 1 to 7. These scores were aggregated for each recruit into a training score feature.

We also included FTO related features for each trainee. Since a recruit has multiple FTOs throughout their training (and each FTO has multiple trainees), we formed a weighted average feature vector across a trainee’s FTOs that includes sex, race, average FTO use of force incidents per year and average awards/recognition received per year. The weighting of the FTO features for each trainee was proportional to the time spent training under each FTO.

We estimated a variance inflation factor (VIF) [3] and Pearson correlation coefficient to identify multicollinearity and removed features with high correlation. Figure 2 displays the correlation matrix after feature selection which ensures that no two features are highly correlated.

4.1 Survival modeling of the time to use of force

To model the survival curve of each officer as a function of their characteristics and the FTO network, we considered two popular methods: Cox proportional hazards and random survival forests.




Figure 2: Correlation matrix after removing features with high correlation.

Cox Regression

The Cox proportional hazards model, also known as a Cox regression, is a semi-parametric model that estimates the relative risk of failure (or hazard) for different groups of individuals based on a set of covariates. Here time t is the time in days since the start of training for each trainee, and failure time is the time of occurrence of a use of force incident involving the trainee.

Let $Z^{(i)} \in \mathbb{R}^d$ be a covariate vector for officer i , the *hazard* for officer i is,

$$\lambda_i(t) = \lambda_0(t) \exp\{\beta^T Z^{(i)}\}$$

and the corresponding survival function is,

$$S_i(t) = S_0(t) \exp\{-\beta^T Z^{(i)}\}$$

where $\beta \in \mathbb{R}^d$ is a vector of trainable parameters. The data for the model is denoted as $D_{\text{surv}} = \{(y_i, \delta_i, Z^{(i)}) : i = 1, 2, \dots, n\}$, where y_i is the minimum of the censoring time C_i (end of observation time) and survival time T_i , and

$$\delta_i = \begin{cases} 1 & \text{if } T_i = y_i \\ 0 & \text{otherwise} \end{cases}.$$

In the above equation, δ_i denotes whether or not an instance is censored.

We train β by maximizing the partial likelihood (under an iid assumption),

$$L(\beta|D_{\text{surv}}) = \prod_{i=1}^n \left[\frac{\lambda_0(t) \exp\{\beta^T Z^{(i)}\}}{\sum_{j \in \Psi(t)} \lambda_0(t) \exp\{\beta^T Z^{(j)}\}} \right]^{\delta_i} + -$$

where $\Psi(t) = \{i : y_i > t\}$ is the subset of officers who “survive” (no use of force incident) past time t .

Random Survival Forest (RSF)

Whereas Cox regression is a parametric, log-linear survival model, random survival forests (RSF) are a non-linear, non-parametric approach for analyzing survival data [8]. RSF is a tree-based algorithm that aggregates an ensemble of decision trees, where each tree is trained on a random subset of the data. The final prediction is made by averaging the predictions over the ensemble of trees. The general algorithm is as follows:

- (1) Draw B bootstrap samples.
- (2) Grow a survival tree for each bootstrap sample. At each node select a subset of predictor variables. Among all binary splits across the predictor variables randomly selected at the node, find the best split determined by the maximum survival difference between daughter nodes. Repeat until stopping criteria is met (for example once each node has only k observations).
- (3) Aggregate information from terminal nodes (nodes for which there are no further splits) and calculate a cumulative hazard function (CHF) for each tree. Average the tree CHFs to obtain the ensemble CHF.

The cumulative hazard function for each terminal node h is given by,

$$\hat{H}_h(t) = \sum_{t_{l,h} \leq t} \frac{d_{l,h}}{Y_{l,h}},$$

where T denotes the set of terminal nodes, $d_{l,h}$ is the number of officers with a use of force incident before time $t_{l,h}$ and $Y_{l,h}$ is the number of individuals at risk at time $t_{l,h}$. Letting $Z^{(i)}$ be the covariate vector for officer i as above, the cumulative hazard function for each tree is then given by:

$$H(t|Z^{(i)}) = \hat{H}_h(t), \quad \text{if } Z^{(i)} \in h.$$

4.2 A 2-opt annealing algorithm for optimizing the survival network

Our goal is to rewire the training network to maximize the expected time to the next use of force incident, averaged over all trainees in the network. While we cannot change individual level characteristics of the officers in training, we can rearrange the FTOs and hence change the network features of each trainee. For example, suppose trainee A has 2 FTOs, C and D, both of whom have a low number of use of force incidents in their histories. Trainee B, on the other hand, has 2 FTOs, E and F, who have a high number of use of force incidents in their histories. Our hypothesis is that by switching C and E, we can improve the expected time of the next use of force incident for trainee B, while also improving the average expected time across A and B. At the same time, we want to maintain the workloads of FTOs, and we may also want to encourage diversity of each trainees FTO group. Therefore we consider a cost of the form,

$$C = \text{mean survival time (MST)} + \lambda * \text{FTO diversity fraction},$$

where the first term is expected time to the next use of force incident, averaged over all trainees, and the second term is the fraction of trainees that have at least one FTO who is not a white male (and λ is a parameter that can control the balance of the two terms).

We then apply 2-opt annealing, where at each iteration we randomly exchange two FTOs between two random recruits. We then keep the new network configuration with a probability that increases

with the improvement of the above cost function (that we want to maximize). Pseudocode for the annealing heuristic is given in the Algorithm 1.

Algorithm 1. Survival Network Annealing

```

1: Initialize network and temperature  $T_0$ .
2: Calculate initial cost (average failure time) using random survival forest,  $C_0$ .
3: for  $i \leftarrow 1$  to  $N$  do
4:   Generate temperature,  $T_i = \frac{T_{i-1}}{i}$ 
5:   Exchange FTO between two random recruit nodes
6:   Update recruit features
7:   Calculate the new cost,  $C_i$ 
8:   Calculate acceptance probability,
       
$$P = \begin{cases} 1 & \text{if } \delta > 0 \\ e^{\frac{\delta}{T_i}} & \text{otherwise} \end{cases}$$

       where  $\delta = C_i - C_{i-1}$ 
9:   if  $\text{rand}() > P$  then
10:     Reject the new network.
11:   else
12:     Accept the new network.
13:   end if
14: end for
```

To initialize the network before annealing, we consider three configurations:

- (1) **Random configuration.** As the name suggests, we started with a random graph with each recruit (trainee) node having 3 random FTOs.
- (2) **IMPD configuration.** In this configuration, we start with the IMPD FTO-trainee graph from 2018-2021.
- (3) **Greedy configuration.** In this configuration, we greedily add FTOs to the trainee network to yield the greatest incremental increase in the expected survival time with each additional FTO.
- (4) **Removing bad apples.** For each of the three above configurations, we also explore the effect of “removing bad apples”, e.g. removing the top 10% of FTOs from the initial network with the highest number of use of force incidents in their history.

In our simulations, there are on average 3 FTOs per recruit and 5-7 recruits per FTO. The annealing algorithm typically converged (no further improvement in the cost) within 35,000 to 100,000 iterations.

5 EXPERIMENTS AND RESULTS

First, we use concordance (c-index) to compare the Cox regression model with the random survival forest. Table 2 reports the scores from the Cox regression and random survival forest on a 75%-25% split of the 189 recruits present in our dataset. Here we find that the RSF outperforms the Cox regression across the three risk sets. We also evaluate the survival models by splitting the test data into quintiles ranked from low to high based on the average forecasted time to use of force across each quintile. Figure 3 presents the actual vs predicted survival time in each group. Due to the superior




Figure 3: Actual vs Predicted mean failure time for 5 officer cohorts in test-set across during training, after 6 months and after 12 month of training risk sets respectively.

performance of the random survival forest, the remainder of our analyses make use of the RSF.

Table 2: Survival Results Summary. Train/test split of 141 and 48 officers, respectively. C-Index shown for models using different risk sets. Random survival model performed best across all risk sets.

Risk Set	Training set	Cox regression (c-index)	RSF (c-index)
During training	train	0.65	0.71
	test	0.57	0.66
After 6 months	train	0.68	0.72
	test	0.63	0.65
After 12 months	train	0.68	0.72
	test	0.64	0.65

In Figure 4, we show the Kaplan-Meier survival curves for officers in different risk sets, disaggregated by the average number of FTO use of force incidents. Here we observe a significant change in the survival probability as FTO use of force history changes. Trainees tend to have a longer time to their first use of force incident when their FTOs have fewer use of force incidents in their history. This result is consistent across training, 6 months and 12 months after training. The mean time to use of force across the three time periods is depicted in the Figure 5.

In Table 3, we show the feature importance along with uncertainty estimates for each feature defined in our model. Here we observe that FTO weighted use of force cases per year is the variable with highest score, and this is true across the three risk sets. We note that the features in Table 3 are the remaining features after removing co-linear features from the model.

We performed an additional experiment where we include second degree trainee-trainee peer effects in the model. Although it points in the direction of a positive effect, this feature was non-significant.

Next we investigate the improvement of the average survival time of the network after applying 2-opt annealing. In Table 4, we show the starting cost C_0 and final cost C_N after annealing for different choices of the initial network, and with or without removal of the top 10% of FTOs with the highest number of historical use of force incidents per year. In the actual IMPD network, the average forecasted time to the first use of force incident for trainees is 64.912 ± 0.763 days. When we optimize the network through annealing (starting from the IMPD configuration), we improve the average

Table 3: Random Survival forest Model Summary. Table displays feature importance score mean along with standard deviation based on c-index attributed to the features defined in our model (difference between c-index with and without feature).

Features	Training		6 months		12 months	
	mean	std	mean	std	mean	std
Overall scores	0.010	0.006	0.002	0.011	0.005	0.012
FTO weighted male	0.027	0.010	0.007	0.006	0.010	0.007
FTO weighted black	0.008	0.003	-0.004	0.005	-0.002	0.004
FTO weighted cases	0.122	0.024	0.075	0.019	0.074	0.020
FTO weighted awards	0.015	0.004	0.001	0.005	0.003	0.006
location north	0.001	0.001	-0.001	0.005	-0.001	0.002
location east	0.002	0.001	0.007	0.008	0.006	0.006
location northwest	0.002	0.001	-0.001	0.001	0.001	0.001
location southwest	0.000	0.001	0.000	0.000	0.000	0.000
Sex male	0.002	0.002	0.004	0.003	0.006	0.004
Race white	0.000	0.003	0.000	0.001	0.001	0.001

time-to-event to 66.944 ± 0.763 days. However, by starting from the greedily constructed initial network, we can improve the average time to 69.405 ± 0.681 days. By removing the 10% of FTOs with the highest number of historical use of force incidents, we can improve the average time further, to 70.410 days after the start of training (an improvement over the original IMPD network of 8.5%). These results are consistent across the three time periods we considered. We also note that a greater improvement is achieved by rewiring the network, rather than removing the top 10% of FTOs. This is consistent with recent research that suggests that removing the top 10% of officers in a network (according to risk) may only yield a 4-6% reduction in incidents [1].

Next we investigate the impact of adding the diversity constraint to the cost function during annealing. In Table 5, we display the mean survival time after annealing, along with the fraction of trainees in the network that have at least one non-white FTO. We observe that the incorporation of a diversity constraint into the annealing cost function leads to an improvement in the fraction of trainees with at least one trainer who is not white or male. For example, in the original IMPD network, 74.6% of trainees have a diverse set of FTOs (as measured by the diversity constraint). After annealing, this fraction goes up to 98.9% ($\lambda = 0.1$) and 100% ($\lambda = 0.5$). Thus each officer in training can be assigned a diverse set of FTOs in




Figure 4: Kaplan-Meier curve showing survival probability vs days disaggregated by the number of FTO use of force cases per year (weighted average across FTOs) during training, after 6 months of training and after 12 months of training.




Figure 5: Mean survival time for officers in different risk set.

Table 4: Starting cost C_0 and final cost C_N after annealing for different choices of the initial network and with or without removal of the top 10% of FTOs with the highest number of incidents.

Risk Set	Network	Method	C_0	C_N
Training	Random	Standard	62.684 ± 0.189	67.523 ± 0.626
		Removing 10%	63.356 ± 0.304	66.907 ± 0.588
	Greedy	Standard	65.840 ± 0.437	69.405 ± 0.681
		Removing 10%	66.316 ± 0.546	70.410 ± 0.726
	IMPD	Standard	64.912 ± 0.585	66.944 ± 0.763
		Removing 10%	65.049 ± 0.602	67.637 ± 0.763
6 months	Random	Standard	90.493 ± 0.467	95.510 ± 0.844
		Removing 10%	91.456 ± 0.542	96.056 ± 0.791
	Greedy	Standard	93.577 ± 0.625	97.712 ± 0.945
		Removing 10%	94.081 ± 0.580	98.571 ± 0.944
	IMPD	Standard	92.619 ± 0.838	96.036 ± 1.098
		Removing 10%	92.772 ± 0.789	96.496 ± 1.015
12 months	Random	Standard	129.917 ± 0.781	141.931 ± 1.690
		Removing 10%	130.166 ± 0.816	139.479 ± 1.638
	Greedy	Standard	136.297 ± 1.459	147.454 ± 2.086
		Removing 10%	136.808 ± 1.318	148.328 ± 2.106
	IMPD	Standard	134.970 ± 1.834	142.959 ± 2.381
		Removing 10%	135.241 ± 1.661	143.880 ± 2.403

the rewired network, while still maximizing the average forecasted survival time.

In Figure 6, we investigate how the survival time depends on the number of historical FTO cases. We consider the existing IMPD

network and vary the average number of FTO cases per year within the range of 0 to 10. We observe a significant decline in the mean survival time of the network, followed by a plateau once the number of FTO incidents per year reaches 4. This is consistent across the three risk sets.

In Figure 7, we show a subgraph of the larger network for 10 recruits before and after annealing. Initially, on the left, the recruits have shorter expected time-to-event, depicted by the smaller node sizes, and the assigned FTOs have a greater number of use of force incidents (indicated by larger FTO nodes). After annealing, on the right, the same recruits are reassigned to different FTOs with fewer use of force incidents in their histories, leading to an increase in the expected time to use of force (reflected in the larger node size).

Finally, we investigate how network rewiring shifts the distribution of time to use of force in Figure 8. We divide recruits into 5 categories based on their expected time to event, e.g. the groups are sorted from highest risk to lowest risk. After annealing, we see that the number of officers in the highest-risk category is significantly reduced, especially for the training time period, whereas we see the greatest increase is in the lowest-risk category (which changes from 9% of recruits to 37% after annealing). Similar patterns are observed after 6 and 12 months. For example, after 12 months the highest risk category goes from having 21% of recruits to 12%, whereas the lowest risk category goes from having 10% of officers to 35%.


6 DISCUSSION

We analyzed a novel dataset from the Indianapolis Metropolitan Police Department on officer field training, subsequent use of force, and the role of network effects from field training officers. We developed a network survival model for the expected time to use of force incidents of police trainees. The model includes individual characteristics of the trainees, as well as network features such as FTO use of force history. We showed that FTO use of force history is the best predictor of trainee time to use of force in the survival model, and that the forecasted average time to use of force across the network can be increased by 8% through rewiring.

Explanations of police brutality are broadly classified into two theories – “bad apples” and “bad institutions”. “Bad apple” theories posit that misconduct is caused by individual factors such as an officer’s age, race, gender, education, and experience. It is believed that these factors lead to deviant behavior among a small group of officers. On the other hand, the “bad institutions” theory focuses

Table 5: Annealing results with diversity constraint added to the cost for several choice of lambda. The initial configuration is the IMPD network.

Risk set	Method	$\lambda = 0.01$			$\lambda = 0.1$			$\lambda = 0.3$			$\lambda = 0.5$			$\lambda = 0.8$		
		Initial frac	Final frac	Final MST	Initial frac	Final frac	Final MST	Initial frac	Final frac	Final MST	Initial frac	Final frac	Final MST	Initial frac	Final frac	Final MST
Training	Standard	0.746	0.979	68.226	0.746	0.989	68.08	0.746	0.995	68.05	0.746	1	67.52	0.746	1	67.596
	Removing 10%	0.73	0.984	68.867	0.73	0.974	69.376	0.73	0.979	68.894	0.73	0.968	69.176	0.73	0.989	68.497
6 months	Standard	0.746	0.974	97.355	0.746	0.989	96.888	0.746	0.979	96.838	0.746	1	96.314	0.746	1	96.451
	Removing 10%	0.73	0.979	97.979	0.73	0.979	98.148	0.73	0.968	98.069	0.73	0.984	97.748	0.73	0.968	98.265
12 months	Standard	0.746	0.968	146.37	0.746	0.963	146.731	0.746	0.934	146.407	0.746	1	143.795	0.746	1	143.653
	Removing 10%	0.73	0.974	148.31	0.73	0.958	149.003	0.73	0.974	148.252	0.73	0.963	150.088	0.73	0.963	149.006

**Figure 6: Non-linear effect of FTO use of force on expected survival time of IMPD network.**

on the organizational factors of a department, such as its culture, hiring practices, and historical context. This approach argues that misconduct is a result of systemic biases within the department [7, 17, 27].

Our approach to rewiring training networks in the present study is grounded in both organizational and individual-level theories. Removing highest risk FTOs or changing individual characteristics of recruits can increase the expected time to use of force; however, considering the network of the entire organization leads to even greater gains. We caution here that gains achievable by network rewiring are likely to be small if applied on their own. We view algorithmic rewiring as an intervention that should complement other changes to policing in the U.S., such as procedural justice training [28], requiring body worn cameras, prohibiting practices such as chokeholds and no-knock warrants, and reforming qualified immunity. Some of these changes have been proposed in H.R. 1280, the George Floyd Justice in Policing Act.

There are several limitations of the present study. Our results show that the average forecasted time to use of force can be increased in a network survival model fit to historical data. While we have attempted to control for shared environment confounding through the use of spatial covariates (representing the assigned district), in general network contagion estimated through regression models is confounded by environment and/or homophily [24]. For example, trainees with a propensity for use of force may have personality characteristics that lead them to being paired with similar FTOs. Field experiments are needed to demonstrate that the potential gains we observed in our simulations are realizable in actual field training

networks. Another interesting line of future research would be to apply network rewiring algorithms to officer networks beyond field training, for example daily patrol pairings and calls for backup. The latter presents mathematical challenges that would need to be overcome, for example how to balance optimizing use of force risk while minimizing call response time. We also note that alternatives to the survival forest may yield more accurate risk assessments. For example, a graph convolution neural network may be better able to learn features related to contagion across higher degree connections. Furthermore, adding more granular time based features would also be a good direction for future research. While the 2-opt heuristic used to solve the present nonlinear integer programming problem empirically converged to a local minimum, future research should concentrate on analyzing the convergence of algorithms for officer pairings and designing efficient algorithms that can scale to larger networks.

Another challenge is the absence of comprehensive national and state level data on use of force. U.S. police departments are decentralized, and they collect and store data in a variety of formats through a wide range of commercial vendors. Further research is needed to determine what elements of the present study are generalizable to other departments, and what aspects are department specific. Moving forward, it will be crucial for agencies to collect and make publicly available comprehensive, accurate data on police use of force, and for researchers to have access to this data to study and understand its impact [18].




Figure 7: Network subgraph of 10 recruits depicting the change in failure time before and after annealing. The yellow nodes in this representation depict recruit officers, while the lavender nodes represent field training officers. The size of the yellow nodes shows the cost category of failure time, with larger nodes indicating a higher failure time and smaller nodes indicating a lower failure time. The size of the lavender nodes, on the other hand, depicts the average number of use of force incidents, with larger nodes indicating a higher number of incidents and smaller nodes indicating a lower number. Yellow nodes getting bigger in the final subgraph after annealing suggests an increase in failure time after being rewired with smaller lavender nodes.

7 ACKNOWLEDGMENTS

This research was supported by NSF grants SCC-2125319 and ATD-2124313, AFOSR MURI grant FA9550-22-1-0380, and the Indiana University Racial Justice Research Fund. We thank the Indianapolis Metropolitan Police Department for providing the data used in this study. Author GM is a member of the board of directors of Geolitics, a data analytics company that serves law enforcement agencies.

REFERENCES

- [1] Aaron Chalfin and Jacob Kaplan. 2021. How many complaints against police officers can be abated by incapacitating a few “bad apples?”. *Criminology & Public Policy* 20, 2 (2021), 351–370.
- [2] Wei Chen, Yajun Wang, and Siyu Yang. 2009. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. 199–208.
- [3] Trevor A Craney and James G Surles. 2002. Model-dependent variance inflation factor cutoff values. *Quality engineering* 14, 3 (2002), 391–403.
- [4] Francesco Fabbri, Yanhao Wang, Francesco Bonchi, Carlos Castillo, and Michael Mathioudakis. 2022. Rewiring what-to-watch-next recommendations to reduce radicalization pathways. In *Proceedings of the ACM Web Conference 2022*. 2719–2728.
- [5] Ryan M. Getty, John L. Worrall, and Robert G. Morris. 2016. How Far From the Tree Does the Apple Fall? Field Training Officers, Their Trainees, and Allegations of Misconduct. *Crime & Delinquency* 62, 6 (2016), 821–839. <https://doi.org/10.1177/0011128714545829>
- [6] James T Haider. 1990. *Field training police recruits: Developing, improving, and operating a field training program*. Charles C Thomas Publisher.
- [7] Jessica Huff, Michael D White, and Scott H Decker. 2018. Organizational correlates of police deviance: A statewide analysis of misconduct in Arizona, 2000–2011. *Policing: An International Journal* (2018).
- [8] Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, and Michael S Lauer. 2008. Random survival forests. *The annals of applied statistics* 2, 3 (2008), 841–860.
- [9] Jessica Jaglois, Nicholas Bogel-Burroughs, and Mitch Smith. 2023. *Initial Police Report on Tyre Nichols Arrest Is Contradicted by Videos*. <https://www.nytimes.com/2023/01/30/us/tyre-nichols-arrest-videos.html>
- [10] Elias Boutros Khalil, Bistra Dilkina, and Le Song. 2014. Scalable diffusion-aware optimization of network topology. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1226–1235.
- [11] Chris J Kuhlman, Gaurav Tuli, Samarth Swarup, Madhav V Marathe, and SS Ravi. 2013. Blocking simple and complex contagion by edge removal. In *2013 IEEE 13th international conference on data mining*. IEEE, 399–408.
- [12] Huayu Li, Yong Ge, Hengshu Zhu, Hui Xiong, and Hongke Zhao. 2017. Prospecting the career development of talents: A survival analysis perspective. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 917–925.
- [13] Hyeyoung Lim and Hoon Lee. 2015. The effects of supervisor education and training on police use of force. *Criminal Justice Studies: A Critical Journal of Crime, Law and Society* 28 (08 2015). <https://doi.org/10.1080/1478601X.2015.1077831>
- [14] John Lu, Sumati Sridhar, Ritika Pandey, Mohammad Al Hasan, and George Mohler. 2019. Investigate transitions into drug addiction through text mining of Reddit data. In *Proceedings of the 25th ACM SIGKDD International Conference*




Figure 8: Officer network rewiring initial and final frequency distribution for different cost category group from greedy configuration using annealing across the 3 risk set

- on *Knowledge Discovery & Data Mining*. 2367–2375.
- [15] Yao Ma, Suhang Wang, Tyler Derr, Lingfei Wu, and Jiliang Tang. 2021. Graph adversarial attack via rewiring. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1161–1169.
- [16] Michael S McCampbell. 1987. Field training for police officers: The state of the art. (1987).
- [17] James P McElvain and Augustine J Kposowa. 2008. Police officer characteristics and the likelihood of using deadly force. *Criminal justice and behavior* 35, 4 (2008), 505–521.
- [18] Kyle McLean, Seth W Stoughton, and Geoffrey P Alpert. 2022. Police uses of force in the usa: a wealth of theories and a lack of evidence. *Cambridge Journal of Evidence-Based Policing* 6, 3-4 (2022), 87–108.
- [19] Sourav Medya, Arlei Silva, Ambuj Singh, Prithwish Basu, and Ananthram Swami. 2018. Group centrality maximization via network design. In *Proceedings of the 2018 SIAM International Conference on Data Mining*. SIAM, 126–134.
- [20] Richard A. Oppel Jr. and Kim Barker. 2020. New Transcripts Detail Last Moments for George Floyd. *New York Times* July 8 (2020).
- [21] Marie Ouellet, Sadaf Hashimi, Jason Gravel, and Andrew V Papachristos. 2019. Network exposure and excessive use of force: Investigating the social transmission of police misconduct. *Criminology & Public Policy* 18, 3 (2019), 675–704.
- [22] Richard Read. 2020. *Derek Chauvin, officer arrested in George Floyd's death, has a record of shootings and complaints.* <https://www.latimes.com/world-nation/story/2020-05-29/chaumin-shootings-complaints-minneapolis-floyd>
- [23] Greg Ridgeway. 2016. Officer risk factors associated with police shootings: a matched case-control study. *Statistics and Public Policy* 3, 1 (2016), 1–6.
- [24] Cosma Rohilla Shalizi and Andrew C Thomas. 2011. Homophily and contagion are generically confounded in observational social network studies. *Sociological methods & research* 40, 2 (2011), 211–239.
- [25] Ivan Y. Sun. 2003. A Comparison Of Police Field Training Officers' And Nontraining Officers' Conflict Resolution Styles: Controlling Versus Supportive Strategies. *Police Quarterly* 6, 1 (2003), 22–50. <https://doi.org/10.1177/1098611102250573> arXiv:https://doi.org/10.1177/1098611102250573
- [26] Thomas W Valente. 2012. Network interventions. *Science* 337, 6090 (2012), 49–53.
- [27] George Wood, Daria Roithmayr, and Andrew V Papachristos. 2019. The network structure of police misconduct. *Socius* 5 (2019), 2378023119879798.
- [28] George Wood, Tom R Tyler, and Andrew V Papachristos. 2020. Procedural justice training reduces police use of force and complaints against officers. *Proceedings of the National Academy of Sciences* 117, 18 (2020), 9815–9821.
- [29] Jiacheng Wu, Forrest W Crawford, David A Kim, Derek Stafford, and Nicholas A Christakis. 2018. Exposure, hazard, and survival analysis of diffusion on social networks. *Statistics in medicine* 37, 17 (2018), 2561–2585.
- [30] Xiaojian Wu, Daniel Sheldon, and Shlomo Zilberstein. 2016. Optimizing resilience in large scale networks. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [31] Yao Zhang, Arvind Ramanathan, Anil Vullikanti, Laura Pullum, and B Aditya Prakash. 2017. Data-driven immunization. In *2017 IEEE International Conference on Data Mining (ICDM)*. IEEE, 615–624.
- [32] Linda Zhao and Andrew V Papachristos. 2020. Network position and police who shoot. *The ANNALS of the American Academy of Political and Social Science* 687, 1 (2020), 89–112.