

Data-Driven Optimization for Police Districting in South Fulton, Georgia

Shixiang Zhu
shixiang.zhu@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

Alexander W. Bukharin
abukharin3@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

Le Lu
helenlule@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

He Wang
he.wang@isye.gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

Yao Xie
yao.xie@isye.gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

ABSTRACT

We redesign the police patrol beat in South Fulton, Georgia, in collaboration with the South Fulton Police Department (SFPD), using a predictive data-driven optimization approach. Due to rapid urban development and population growth, the existing police beat design done in the 1970s was far from efficient, which leads to low policing efficiency and long 911 call response time. We balance the police workload among different city regions, improve operational efficiency, and reduce 911 call response time by redesigning beat boundaries for the SFPD. We discretize the city into small geographical atoms, which correspond to our decision variables; the decision is to map the atoms into “beats”, the basic unit of the police operation. We first analyze workload and trend in each atom using the rich dataset, including police incidents reports and U.S. census data; We then predict future police workload for each atom using spatial statistical regression models; Lastly, we formulate the optimal beat design as a mixed-integer programming (MIP) program with continuity and compactness constraints on the beats’ shape. The optimization problem is solved using simulated annealing due to its large-scale and non-convex nature. The simulation results suggest that our proposed beat design can reduce workload variance among beats significantly by over 90%.

CCS CONCEPTS

• **Applied computing** → **Multi-criterion optimization and decision-making**; • **Mathematics of computing** → *Combinatorial optimization*.

KEYWORDS

data-driven optimization, predictive policing, police beat re-design

1 INTRODUCTION

The City of South Fulton, Georgia, was recently established in May 2017 from previously unincorporated land outside Atlanta. It is now the third-largest city in Fulton County, Georgia, and serves a population of over 98,000, among which 91.4% are black, or African American [33]. South Fulton is a historic area renowned for its art and activism. Despite this, the city has often faced the challenge of climbing crime rates and long police response times. In a 2019 survey, 46.48% of residents responded that they do not feel safe in South Fulton. In the same year, the South Fulton City Council made




Figure 1: City-wide police district map of South Fulton, GA. There were 7 beats, which was initially designed in the 1970s. The city boundary is highly irregular which requires intricate design of police beats.

it clear that their number one priority was to make South Fulton safer [16].

The South Fulton Police Department (SFPD) is the main policing force in the city. From 2019 to early 2020, our team worked with the SFPD to improve their police operation efficiency. Our project specifically focused on redistricting beat configurations (by completely re-drawing the beat boundaries and changing the number of beats), aiming to rebalance SFPD officers’ *workload* (total amount of working time). The initial analysis identified that workload unbalance among different areas of the city was caused by an outdated beat design that had not been changed for over five decades; the inefficient beat design, in turn, lead to long 911 call response time in some areas.

Previously, the police operation of South Fulton was according to seven police *beats*, which divide the city geographically as shown in Figure 1. 117 police personnel were allocated to the beats for patrolling and responding to the 911 calls [7]. Typically, at each shift, one response unit (usually a police car with one to two officers) answers all the 911 calls that occurred in a certain beat. If the response unit is busy handling another incident, nearby available response units may be dispatched by the operator to answer the call.




Figure 2: Distribution of 911 calls-for-services requests in South Fulton, GA. Blue shaded area is the city limit of South Fulton. Blue dots are locations of requests. The requests are unevenly distributed among different regions.

The most recent South Fulton police beat redesign occurred in the 1970s – almost five decades ago. Since then, the area (which eventually became the City of South Fulton) has undergone tremendous urban growth that drastically changes its landscape. The U.S. Census Bureau estimated that South Fulton’s population has increased by 13.7% from 2010 to 2018 [33]. The city’s rapid development has led to a significant increase in police workload, exacerbated by the difficulty in officer recruitment and retention faced by the SFPD. Moreover, demographic and traffic pattern changes also create an unbalanced workload among different regions. Figure 2 shows the distribution of 911 calls, recorded by real 911-call reports provided by SFPD from 2018 to 2019. The figure shows some beats faced a significantly higher workload than others. For example, police officers in the city’s southeastern area respond to more calls than those in the western region.

Since the seminal work by R. Larson and others [17, 20], researchers have recognized that beat configuration may significantly impact police response time to 911 calls and operational efficiency. In particular, the area and shape of beats determine the workload and travel time in that beat. Hence, it is critical to design the boundaries of beats to balance the workload.

Outline. We redesigned the police patrol beats in the City of South Fulton using a data-driven optimization approach. The outline of our approach is summarized in Figure 3. Our objective is to balance police workload in each beat by redrawing beat boundaries. First, we divided the geographical areas of the city into a large number of “atoms”. Then, we estimated the workload in each atom using police reports data and census data, including population and socio-economic factors. These steps are described in Sec. 2 and 3. Based on the workload estimation, we developed statistical models to predict police workload in the next few years (Sec. 4). We then formulate the beat redesign problem as a clustering problem: each beat is formed with a cluster of atoms. This clustering problem is formulated and solved using mixed-integer programming (MIP), where the objective function is a metric of workload unbalance (defined as the workload variance across all beats). We also impose constraints that require beats to be contiguous and compact so that they are not irregularly shaped. The problem formulation is described in Sec. 5. To tackle the computational complexity of solving a large-scale optimization problem, we developed a simulated




Figure 3: An illustration for data-driven optimization framework of police beat redesign.

annealing based approach with efficient solution exploration. We also study the districting with the different number of beats and find the optimal number of beats with the highest cost-effectiveness. Numerical results (Sec. 6) show that our proposed beat design can reduce workload variance among different regions by over 90%. In January 2020, together with the SFPD, we presented our final redesign plan to the South Fulton city council, which was officially approved for implementation.

Contribution. Our work proposes a new data-driven framework that integrates data, statistical prediction, and optimization in the context of police beat design. Previous works in the predictive policing literature tend to focus on only the prediction aspect. The operations research literature often studies police zone design based on analyzing stochastic models without explicitly considering data sources. We take advantage of the availability of abundant data and adopt a new data-driven approach: the workload and other important parameters for optimization are estimated and predicted from data. From a methodological perspective, we use geo-spatial atoms to define city boundaries and police beat boundaries. This approach enables accurate workload prediction by correlating historical police data with the census data and beat design optimization.

Our project also had a significant societal impact and directly improved the police operations of the SFPD and the safety of residents in South Fulton. It is worth mentioning that although we focus on the study of police beat redesign in South Fulton, our method can be applied to other cities facing similar issues.

Related work. Police districting (designing beats or zones) is a classical problem studied in operations research dating back to the 1970s (see the seminal work [17] and the surveys by [5, 13] for reviews). [11] is one of the earliest works that study optimal beat allocation using integer programming. [2] considers the beat allocation problem to minimize response time for police service calls. In particular, the paper also considers overlapping beats, where multiple patrol officers share one patrol area. [6, 18] use queueing models to estimate travel time. In particular, our proposed data-driven model includes the travel time in the workload calculation. [19] introduces a heuristic approach to the design of beats with implementation in Boston. [4] considers fairness issues of police zone design. We remark that most classical works rely on analyzing stochastic models for police workload estimation, which usually requires stringent assumptions, e.g., calls arrive according to homogeneous Poisson processes (with the notable exception of [18]). Here, rather than obtained from stochastic models, we take advantage of the availability of abundant data and adopt a data-driven approach: the workload and other important parameters for optimization are estimated and predicted from data.

There is also a large body of works on other types of geographical districting problems, such as political districting. This includes the pioneering work [10] studies political districting using integer

programming. Their method is extended by [31] for other geographical districting problems. A few other works [8, 23–25, 35, 37] apply meta-heuristics (e.g., genetic algorithms, simulated annealing) to geographical districting, which usually lack optimality guarantees. Geographic districting often include criteria such as contiguity [10, 14, 22, 23, 26, 35] and compactness [10, 27, 38], which are also important in the police zone design context. However, political districting has different considerations than police districting.

In the last decade, we have seen the rise of predictive policing, i.e., the use of mathematical and statistical methods in law enforcement to predict future criminal activity based on past data. Its importance has been even recognized by Time magazine that in November 2011 named predictive policing as one of the 50 best inventions of 2011 [15]. The RAND Corporation and the National Institute of Justice of the United States (NIJ) also acknowledge the need for taking a step forward and developing explicit methodologies and tools to take advantage of the information provided by predictive policing models to support decision makers in law enforcement agencies [28]. We remark that most classical works do not leverage the historical operational data and rely on analyzing stochastic models for police workload estimation, which usually requires stringent assumptions, e.g., calls arrive according to homogeneous Poisson processes. Here, rather than obtained from stochastic models, we take advantage of the availability of abundant data and adopt a data-driven approach: the future workload and other essential parameters for optimization are estimated and predicted from data.

2 DATA

We start by describing the various sources of data used for South Fulton police beats reconfiguration, including 911 calls-for-service reports, geographical data of the city, and the socio-economic data collected by the American Community Survey (ACS) from the U.S. Census Bureau.

911 calls-for-service data. The SFPD provides comprehensive 911-call reports between May 2018 to April 2019, which contains 69,170 calls in total (Figure 2). The recorded 911 calls cover more than 600 categories of incidents, including assaults, terrorist threats, domestic violence, robbery, burglary, larcenies, auto-thefts, etc. These reports are generated by mobile patrol units in the city, which handle 911 calls 24/7. Teams of *response units* (police cars and officers) are assigned to patrol city streets, and answer calls for service. When a 911 call for a traffic incident comes in at the *call time*, a new incident record will be created at the dispatch center, and the call location will be recorded. The operator assigns an officer to handle the call. The unit arrives at the scene and starts the investigation. Once the police complete the investigation and clear the incident, the police report will be closed and record the *clear time*. The time interval that it takes police to process the call between the call time and the clear time is called *processing time*. The police workload is calculated using both the geolocation data and 911 call processing time data (The calculation method which will be discussed in more detail in Sec. 3). The geolocation consists of the GPS location of reported incidents. From the geographical data of South Fulton, we are also able to identify which beat each incident is located.

GIS data & beat configuration. Geographic information system (GIS) data contain the geographical information of the city’s and




Figure 4: (a-c): Raw data for demographic factors of South Fulton, GA in 2019, from American Community Survey, organized by census blocks. (d-f): Corresponding atomized census data of South Fulton, GA, in 2019.

beats’ boundaries, which are extracted from Fulton County Special Services District digest parcel data [12]. Geographically, the city boundary of South Fulton is quite irregular with jagged edges, holes, and disconnected segments (Figure 1). This irregularity is due to the formation of the City of South Fulton, with the city being a new combination of all the unincorporated land in southwest Fulton County. Currently, there are seven beats in the City of South Fulton. As shown in figure 1, beat (district) 1, 2, 3, and 4 include larger areas that are relatively compact, while remaining beats contain smaller scattered areas. The irregular shape of the city brings difficulty to police officers while reaching locations of requests and patrolling. Moreover, as the busiest airport globally, the Hartsfield-Jackson Atlanta International Airport is situated east of the city, which significantly adds to the city’s workload disparity.

Census data. The American Community Survey (ACS) collected by the U.S. Census Bureau provides comprehensive information about the population, demographic, and economic status of different Georgia areas. Unlike the census, which takes place every ten years, the ACS is conducted once per year. Some demographic factors are useful in predicting future workload (by correlating the city’s socio-economic profile with the workload). These factors contain essential information about the development and economic growth of the city. Besides, census data is organized by census blocks, as shown in Figure 4 (a - c), which is also different from the geographical atoms we consider for our study. In our discussions with the SFPD, we selected eight most influential factors that play a vital role in determining the police workload, such as population and median rent, school enrollment, and the average year structures were built. The full list of census factors we are considering has been shown in Table 1.

3 DATA PREPROCESSING

In this section, we describe three key steps in data preprocessing before performing the beat design. In particular, we need to address




Figure 5: South Fulton region is partitioned into 1,187 square geographical atoms. Color indicates the beat membership.

the following challenge in using the data: how to align time resolution and spatial resolution from the raw data with what we need in the design.

Geographical atoms. To accurately capture changing demographics and determine the new boundaries for each police beat, we define high-resolution geographical atoms by creating artificial polygons of identical size as our geographical atoms to break up the city. The optimal beat design can be found by aggregating multiple adjacent polygons. The size of geographical atoms is essential to our design’s performance since it determines the number of variables in the optimization and the precision of the workload estimation. There is a trade-off between computational efficiency and model accuracy in determining the size of geographical atoms. If the atoms’ size is too large, then we are unable to capture community demographics accurately; if the size of the atoms is too small, then the problem will become computationally intractable. After rounds of discussion with the SFPD, we decide atoms to be a square area with a side length of 0.345 miles, roughly the city block size. This allows us to estimate the local workload accurately while resulting in a reasonable number of decision variables in our optimization problem. The atomized map of the city was generated by intersecting the city boundary with a grid of atoms, resulting in a new grid of 1,187 geographical atoms, as shown in Figure 5. The police workload estimation and prediction will be performed based on these predefined geographical atoms. Formally, let $i \in \mathcal{I} = \{1, \dots, I\}$ denote the i -th atom and $k \in \mathcal{K} = \{1, \dots, K\}$ denote the k -th beat in our design. Let the binary decision variable $d_{ik} \in \{0, 1\}$ denote whether or not atom i is assigned to beat k . A particular beat design is a unique graph partition determined by a matrix $D = \{d_{ik}\} \in \{0, 1\}^{I \times K}$. For each i , it satisfies $\sum_{k=1}^K d_{ik} = 1$. Given the beat design D , the set of atoms assigned to beat k is denoted by $\mathcal{S}_k(D) = \{i : d_{ik} = 1\} \subseteq \mathcal{I}$. Figure 5 also shows the discretization of the existing beat configuration, where atoms with the same color represent a police beat.

Census data atomization. A major challenge for estimating the socio-economic data for each geographical atom using census data is the inconsistency between census blocks and geographical atoms,

where, as shown in Figure 4 (a-c), census blocks usually have a much larger area than geographical atoms. Here we need to perform a spatial interpolation to align the census data with our geographical atoms. Specifically, we assume the census data, such as population, in the same geographical atom, are evenly distributed. The data of each census factor in a geographical atom can be estimated by proportionally dividing the value in the census block where the atom falls into. The weight of the portion that an atom takes from a census block can be measured by the proportion between their areas. As shown in Figure 4 (d - f), the census data collected by census blocks have been discretized into geographical atoms. Given historical census data in the month $\ell \in [L - L_0, L]$, where L and L_0 denote the last month and the time span of the historical data, respectively. The preprocessed census data is denoted as a tensor $X = \{x_{i\ell m}\} \in \mathbb{R}^{I \times L_0 \times M}$, where each entry $x_{i\ell m}$ indicates the value of the census factor $m \in \mathcal{M} = \{1, \dots, M\}$ in atom i and month ℓ .

911 calls-for-service data preprocessing. We estimate the police workload for each geographical atom using the 911 calls-for-service dataset. The workload of each 911 call is evaluated by its processing time, i.e., the total time that the police spend on traveling and the investigation. We calculate the workload by two steps: (1) count the number of 911 calls occurred in the i -th atom in ℓ -th month, denoted as $N_{i\ell}$; (2) estimate the total workload for the i -th atom in the ℓ -th month by multiplying $N_{i\ell}$ by the average processing time, denoted as $w_{i\ell}$. The count of 911 calls will be further used as the predictor in our spatial regression model, which will be discussed in Sec. 4.

4 POLICE WORKLOAD PREDICTION

Predicting the police workload is particularly challenging. Although we assume the call arrival rates are time-homogeneous, we observe in the actual data that the call arrival rates have a significant seasonality pattern and yearly trend, as well as correlation over adjacent geographical areas. Therefore, we propose a spatio-temporal model to predict future call arrival rates. We assume that the call arrival rate $\lambda_{i\ell}$ for atom i in a month ℓ is a constant. Thus, each beat is a homogeneous Poisson process with rate $\lambda_{i\ell}$. The arrival rates $\Lambda = \{\lambda_{i\ell}\} \in \mathbb{R}_+^{I \times L_0}$ can be approximated by $N_{i\ell}$, where $L_0 = 12$. We learned from the SFPD that, the occurrence of 911 calls is highly correlated with population and economic status of the beat and its neighborhood. We predict the arrival rate $\lambda_{i\ell}$ in the future month $\ell = L + t, t = 1, 2, \dots$ using a linear model that regresses the arrival rate to other endogenous variables (arrival rates in other beats) and exogenous factors (demographic factors). As shown in Table 1, we consider $M = 8$ demographic factors, which are statistically verified to be good predictors, including population, education level, and household income. Specifically, we use the spatially lagged endogenous regressors [30] defined as

$$\lambda_{i\ell} = \sum_{(i,j) \in \mathcal{A}} \alpha_{ij} \lambda_{j\ell} + \beta_0 \lambda_{i,\ell-1} + \sum_{t=1}^p \beta_t^T X_{i,\ell-t} + \epsilon_i, \quad \forall \ell \in [L - L_0, L],$$

where p is the total number of past months of data that we consider for fitting the regressor, which in our case was 1. The adjacency matrix $A = \{\alpha_{ij}\} \in \mathbb{R}^{I \times I}$ specifies adjacency relationships between atoms. The temporal coefficient $\beta_0 \in \mathbb{R}$ specifies the influence of the last month. The coefficient $\beta_t \in \mathbb{R}^M, \forall 1 \leq t \leq p$ specifies

Table 1: Variables used for workload prediction

PREDICTOR	REGRESSION COEFFICIENT	P-VALUE
POPULATION	439.558	0.007
NUMBER OF HOUSING UNITS	158.440	0.019
SCHOOL ENROLLMENT	79.236	0.008
MEDIAN HOUSEHOLD INCOME	59.420	0.000
MEDIAN NUMBER OF ROOMS	-10.560	0.006
MEDIAN AGE	-7.421	0.001
MEDIAN HOUSE PRICE	-16620	0.000
AVERAGE YEAR BUILT	170.140	0.003

correlations with census factors and error term ϵ_i are spatially correlated. The set of adjacency pairs is defined by $\mathcal{A} = \{(i, j) : i, j \text{ are adjacent in } \mathcal{G}; i, j \in \mathcal{I}\}$. The graph $\mathcal{G}$ is given by associating a node with every atom and connecting two nodes by an edge whenever the corresponding atoms are geographically adjacent. Here, we capture the spatial correlation between data using the standard spatial statistics approach, by assuming ϵ_i to be spatially correlated with correlation depending distance between two locations [29].

5 BEAT REDESIGN OPTIMIZATION

In this section, we introduce our objective and solution methods to the beat redesign optimization problem. We develop an optimization framework to shift beat boundaries, where artificial geographical atoms were assigned to beats while balancing the workload. We formulate this problem as minimizing the workload variance by re-configuring the beat plan with constraints, including the continuity and compactness of beats.

5.1 Objective

Our goal is to shift beat boundaries and make inter-beat workload distribution even. Based on the discussion with the police, we choose the objective function as the workload variance among different zones, which quantitatively measure the police workload imbalance between zones from a macro view. The objective of this problem can be formulated as minimizing the inter-beat workload variance $Z(D)$ given a beat design D :

$$\begin{aligned}
 & \underset{D}{\text{minimize}} \quad Z(D) := \sum_{k=1}^K \left(w_{k\ell}(D) - \frac{\sum_{k=1}^K w_{k\ell}(D)}{K} \right)^2 \\
 & \text{subject to} \quad \sum_{k=1}^K d_{ik} = 1, \quad \forall i \\
 & \quad \text{contiguity and compactness for each beat.}
 \end{aligned} \tag{1}$$

Recall that the matrix $D = \{d_{ik}\} \in \{0, 1\}^{I \times K}$ represents decision variables, where binary variable $d_{ik} \in \{0, 1\}$ indicates whether or not geographical atom i is assigned to beat k ; and $w_{k\ell}(D) = \sum_{i \in \mathcal{I}_k(D)} w_{i\ell}$ represents the total workload in beat $k = 1, \dots, K$ in month ℓ . The variance is a quadratic function of the workload in each beat, which implies that the objective function is convex with respect to the decision variables. A smaller variance indicates a more balanced inter-beat police workload. The constraints will be explicitly defined in Sec. 5.2.

5.2 Compactness and contiguity constraints

In addition to balancing the police workload, it is desirable that the beat shapes are contiguous and compact. In fact, the police never used a quantitative measure of compactness to declare the plans unsuitable. Instead, the police have simply disallowed plans with long and thin or snakelike districts. In other words it appears that the police have evaluated compactness only visually. Since it is not obvious how to determine an acceptable compact design, we choose to minimize the workload variance based on the discussion with the police; but it should be understood that compactness is in reality a loose constraint rather than an objective. Therefore, we formulate the *contiguity* and *compactness* criteria as a set of linear constraints [10, 27, 32, 38] by introducing additional variables: f_{ijk} is the flow from atom i to atom j in beat k ; h_{ik} equals to 1 if atom $i \in \mathcal{I}$ is selected as a sink in beat $k \in \mathcal{K}$, otherwise 0; q is the maximum beat capacity. Hence, there are 21,170,145 variables with 63,421,410 constraints in total.

Contiguity constraints. Contiguity constraints are imposed on each beat using the flow method [32]. For each beat k , there is a flow f_{ijk} on the graph, where f_{ijk} denotes flow from i to j . Each beat has a hub vertex whose net flow is at most the number of vertices in the beat, less one. Each other vertex in the beat has a net flow of at most -1 . This ensures that there is a path of positive flow from any vertex in the beat to the hub, implying contiguity.

Specifically, constraints (2a) represent the net outflow from each beat. The two terms on the left indicate, respectively, the total outflow and total inflow of atom i . If atom i is included in beat k but is not a sink, then we have $d_{ik} = 1, h_{ik} = 0$, and thus atom i must have supply ≥ 1 . If atom i is included in beat k and is a sink, then we have $d_{ik} = 1, h_{ik} = 1$, and thus atom i can have demand (negative net outflow) $\leq q - 1$. If atom i is not included in beat k and is not a sink, then we have $d_{ik} = 0, h_{ik} = 0$, and thus atom i must have supply 0. If atom i is not included in beat k but is a sink, then we have $d_{ik} = 0, h_{ik} = 1$, and the rest of $d_{\cdot k}$ are forced to be 0, that is, no atoms are selected. Constraints (2b) specify the number of atoms that can be used as sinks. Constraints (2c) ensure that each beat must have only one sink. Constraints (2d) ensure that there is no flow into any atom i from outside of beat k (where $d_{ik} = 0$), and that the total inflow of any atom in beat k (where $d_{ik} = 1$) does not exceed $q - 1$. Constraints (2e) make sure unless a atom i is included in beat k , the atom i cannot be a sink in beat k . Constraints (2f) and (2g) ensure that there are no flows (inflows and outflows) between different beats which forces eligible contiguity.

$$\sum_{(i,j) \in \mathcal{A}} f_{ijk} - \sum_{(i,j) \in \mathcal{A}} f_{jik} \geq d_{ik} - qh_{ik}, \quad \forall i, k, \tag{2a}$$

$$\sum_k \sum_i h_{ik} = K, \tag{2b}$$

$$\sum_i h_{ik} = 1, \quad \forall k, \tag{2c}$$

$$\sum_{(i,j) \in \mathcal{A}} f_{jik} \leq (q-1)d_{ik}, \quad \forall k, \tag{2d}$$

$$h_{ik} - d_{ik} \leq 0, \quad \forall i, k, \tag{2e}$$

$$f_{ijk} + f_{jik} \leq (q-1)d_{ik}, \quad \forall i, k, \tag{2f}$$

$$f_{ijk} + f_{jik} \leq (q-1)d_{jk}, \quad \forall j, k, \quad (2g)$$

$$d_{ik}, h_{ik} \in \{0, 1\}, \quad \forall i, k, \quad (2h)$$

$$f_{ijk} \geq 0, \quad \forall i, j, k, \quad (2i)$$

Compactness constraints. Compactness is defined as geographical compactness with distance compactness and shape compactness [27, 38]. For distance compactness, a district is feasible only if the distance between population units must be less than a specified upper bound. For shape compactness, a district is feasible only if the square of the distance's maximum diameter divided by the district's area must be less than another upper bound [10].

Following the existing literatures, we add two additional linear constraints (3a), (3b) to ensure the compactness of beats. For each atom i , let A_i be the area of i , and for each pair of atoms i and j , let l_{ij} be the square of the distance between the centroids of the beats. We also have a parameter $c_1, c_2 > 0$ controlling the degree of compactness.

$$l_{ij}e_{ijk} \leq c_1, \quad \forall i, j, k, \quad (3a)$$

$$l_{ij}e_{ijk} \leq c_2 \sum_{i=1}^K d_{ik}A_i, \quad \forall i, j, k, \quad (3b)$$

5.3 Heuristic approximation

Three methods were discussed in our experiments to search for optimal police beat design. The greedy algorithm serves to generate new beats iteratively and confirms the optimal number of beats for the future redesign. Following the greedy redesign, we adopt a heuristic optimization approach to find the beat design in contrast to the mixed-integer programming (MIP) approach.

Greedy search. To determine the optimal number of beats in the final design, we perform an iterative greedy algorithm, which attempts to generate new beat greedily for the design for each iteration while preserving the original structure of the existing beat as much as possible. Intuitively, more beats may result in a more balanced workload distribution. However, the manpower of the SFPD and resources of the South Fulton City Council are limited. It is unrealistic to deploy such a design with a large number of beats. Hence, we adopt the Greedy algorithm to explore the optimal number of beats in our design. The procedure for “Greedy” creating new beat designs is demonstrated as follows.

For the n -th iteration, we define D_n as the beat design, and K is the number of beats at the last iteration. For the predicted workload in month ℓ , the greedy algorithm can be performed by selecting the beat k in D_n with the largest workload, i.e., $\arg \max_k \{w_{k\ell}(D_n)\}_{k \in \mathcal{K}}$. Then we split up the beat k evenly into two beats using the K-means algorithm, where each atom in the beat is considered as a point. This will lead to generating a new beat, i.e., $K := K+1$ and $\mathcal{K} := \mathcal{K} \cup K$. The above process can be carried out iteratively until we find the design with the optimal number of beats.

We visualize our greedy design with different number of beats in Figures 8. As seen from the result, the beat with the highest workload, shown in red, is split in each iteration as a result. We also examine the variance of beat workload versus number of beats, and

find the optimal number of beats, which will be further discussed in Sec. 6.

Mixed-integer programming. Mathematical programming models are essential tools for modeling and solving redistricting problems, which can guarantee the optimality of the obtained solutions, are mostly based on mixed-integer programming (MIP). However, as shown in Sec. 5.2, the problem involves a large number of variables, including 21,134,535 continuous variables and 35,610 binary variables, as well as a set of additional linear constraints needed to be satisfied. In practice, the problem itself of searching for the global optimal design is computationally intractable and hard to be implemented on a large scale.

Heuristic search. A metaheuristic method, simulated annealing (see, e.g., [3]), has been widely adopted in solving the large-scale combinatorial optimization problem. The simulated annealing algorithm explores the neighborhood of the current solution and decides a better substitution randomly. Simulated annealing can achieve reasonable performance in practice for various settings, although there are very limited theoretical performance guarantees [1, 21, 34]. In particular, in our setting, we use the current/existing partition as an initial solution. Based on this, a new solution can be founded by selecting from a set of candidate solutions. The set of candidate solutions is typically constructed as “neighboring” solutions to the current solution without breaking contiguity.

Specifically, in the n -th iteration, our simulated annealing algorithm performs the following acceptance-rejection sampling. Suppose the starting partition is $\mathcal{P}_n$. For instance, we can take the existing partition as an initialization. The next partition $\mathcal{P}_{n+1}$ is selected from a set of candidate partitions defined as $\mathcal{S}_{n+1}$ and $\mathcal{P}_{n+1} \in \mathcal{S}_{n+1}$. The candidate partitions in $\mathcal{S}_{n+1}$ satisfy contiguity and balance constraints. We randomly choose one of these candidate partitions $\mathcal{P}_{n+1} \in \mathcal{S}_{n+1}$, and evaluate a score

$$P(\mathcal{P}_{n+1}, \mathcal{P}_n | T) = \begin{cases} 1, & Z(\mathcal{P}_{n+1}) < Z(\mathcal{P}_n), \\ \exp\{|Z(\mathcal{P}_{n+1}) - Z(\mathcal{P}_n)|/T\}, & \text{otherwise.} \end{cases}$$

where $Z(\cdot)$ denotes the cost associated with a partition (e.g., the compactness shown in (2)), T is a pre-specified temperature parameter that determines the speed of convergence, and the P is the acceptance probability. We generate an independent uniform random variable $U \in [0, 1]$. The proposed partition is accepted if $P(\mathcal{P}_{n+1}, \mathcal{P}_n | T) \geq U$. We refer to an update of the proposed partition as a *transition*. Note that there is a chance that the transition happens from a “low-cost” partition to a “high-cost” partition, and this “perturbation” will prevent the algorithm from being trapped at a local sub-optimal solution. The choice of the set candidate partitions $\mathcal{S}_{n+1}$ is critical for the performance of simulated annealing, which involves the trade-off between exploration and exploitation. Below, we introduce two strategies for candidate partitions and explore two types of “neighbor” partitions based on square and hexagonal grids, respectively.

As illustrated in Figure 6, we first consider the following simple heuristic in constructing the candidate set. This allows us to search for local optimal partitions at a reasonable computational cost [36]. The candidate set contains all partitions that swap a single vertex assignment at the boundary of the current partition. This simple heuristic is easy to implement since the number of such candidate





Figure 6: Illustrations of two approaches for candidate partitions based on one-swapping neighborhoods. Red and blue boxes represent vertices in different parts. The thick black line represents the boundary of two parts. The left panel shows the current partitions $\mathcal{P}_n$; the middle and the right panels show the candidate partition sets for the next iteration $\mathcal{S}_{n+1}$.



(a) 2020 prediction (b) 2021 prediction (c) 2022 prediction

Figure 7: Workload prediction where dark lines outline boundaries of beats and the color depth represents the level of the atom workload in each year.

partitions is usually small (because we only swap one end of the boundary edge). However, on the other hand, such candidate sets may contain partitions that are still too similar to the current partition. Therefore, we will consider the following alternative strategy.

6 RESULTS

In this section, we present our numerical results and final beat redesign for the City of South Fulton.

Workload analysis and prediction. The most important metric for evaluating imbalance we considered is workload variance over beats. As we defined in Sec. 5, the variance is the sum of the squared deviation of the beat workload from its mean. To fully understand the workload imbalance situation, it is necessary to show how the existing configuration exacerbates the unbalance of workload over beats in the past and how the existing configuration will impact the future. Figures 7 summarizes the predicted workload distribution over the entire city for the next four years from 2020 to 2022. As we can see from the map, there is a clear trend that the general workload level continues to increase, and the major workload concentrates on particular areas (such as College Park in the east of the city and I-285 & I-20 in beat 4). Due to the increasing growth of South Fulton and urban sprawl, this trend is leading to a police workload imbalance.

Optimal beat number. When creating a beat design, the most important metric for evaluating imbalance is the workload variance over beats. However, for determining the optimal number of beats in the design, we also need to consider the cost associated with adding more beats, which includes the cost of additional training, hiring

new officers, and so on. Therefore, there is a trade-off to minimize the workload variance while avoiding unnecessary costs for adding new beats. Figure 8 presents comparisons between existing beat design, designs generated by greedy exploration algorithm, and the proposed design. Figure 9 shows that as we first begin to increase the number of beats, the workload variance decreases sharply before 15 beats. We have shown that there are diminishing returns as we further increase the number of beats beyond 15. Therefore, we call 15 the optimal number of beats and the corresponding 15-beat greedy design will be used as an initialization of the simulated annealing for further refinement.

Proposed beat design. The initial report in 2019 contained beat-wise workload prediction for the next three years (2020, 2021 and 2022), and proposed three candidate designs with similar beat shifts that all attains the best workload balance. In Table 2, we list the predicted annual workload in each beat, total workload, and workload variance. After the plans were reported to the police, we met several times to deliberate the various trade-offs, and held police-engagement meetings to elicit feedback from the patrol force. Police Deputy Chief and couples of key senior officers also participated in these discussions and voiced comments. The new design and the previous existing design have been both presented in Figure 8. This design is preferred by the police for three major reasons: (2) this plan makes the minimal changes based on the existing police zone configuration in comparison with other candidate plans, which minimizes the implementation cost in practice; (2) the workload variance has been drastically reduced by 89% ~ 92% by increasing the number of beats to 15; (3) the proposed plan achieves a lower level of workload variance as well as a smaller variance increment in the future year 2021.

Staffing level analysis. We quantify our potential police response workload by converting the workload in each beat into hours per day. Table 2 shows real workload distribution in 2019 and predicted workload distribution in 2021 under different designs, respectively. Entries of the table suggest the number of hours per day, a police officer would expect to be responding to 911 calls. As we can see, our proposed beat design drastically reduces the beatwise workload. In particular, the proposed design results in a decrease in workload

Table 2: Summarization of workload per beat.

BEAT NUMBER	WORKLOAD IN 2019 (HOURS/DAY)			WORKLOAD IN 2021 (HOURS/DAY)	
	EXISTING	GREEDY	REFINED	GREEDY	REFINED
1	38.59	17.15	17.15	18.05	18.05
2	24.84	24.84	23.56	27.09	25.61
3	32.84	18.78	20.08	17.91	19.91
4	34.44	17.45	17.08	16.83	16.14
5	65.94	22.10	20.31	21.40	19.32
6	38.44	14.69	18.30	14.54	16.73
7	34.96	17.55	19.99	17.67	20.01
8	N/A	12.51	12.51	11.66	11.66
9	N/A	10.79	10.79	11.10	11.10
10	N/A	21.45	21.87	21.45	21.87
11	N/A	23.75	19.33	22.2	22.62
12	N/A	17.41	17.41	23.40	21.60
13	N/A	17.00	16.82	16.70	15.87
14	N/A	20.53	18.89	19.99	17.81
15	N/A	14.06	15.94	13.18	15.93
VARIANCE	142.91	15.12	10.13	18.269	13.15




Figure 8: Greedy beat designs where dark lines outline boundaries of beats and the color depth represents the level of the beat workload. The scale is adjusted in each image.

variance of over 85% comparing to the existing design, making policing more equitable in the city.

In the City of South Fulton council meeting, the city council emphasizes the importance of community engagement from the police force. Thanks to our beat design, the police workload per day in each beat can be reduced drastically; this will allow police officers to participate in community events and start pro-active patrols. This is a huge difference from the past 50 years, where police officers have been going from call to call on their entire shift. Additionally, the staffing level prediction gives the SFPD how many officers they need to handle the 911 calls in a beat. They then can recruit more officers for the sole purpose of community engagement and pro-active patrolling if they desire.

7 IMPLEMENTATION

In January 2020, we submitted the final report to the South Fulton Police Department and the South Fulton City Council. The report was reviewed by Police Chief Meadows, Deputy Police Chief Rogers, and Mayor Bill Edwards. Our report analyzed the police workload and proposed a detailed redistricting plan. Our redistricting plan

mainly changed in four areas (Figure 8l): We add three new beats in the southeast of the city near College Park, the area with the highest workload. The biggest beat in the west of the city is split into two beats. We add a beat in the north of the city near the airport. The southern beat is also split into two. In total, the redistricting plan has reduced the response time throughout the city and rebalanced the police workload between the fifteen beats.

Later that month, the South Fulton City Council approved the new beat design. The South Fulton Police Department plan to implement the new beat design in early 2020. The new beat design was praised by the city council, as some council members said that our beat design and study has been long needed and that it sets an example for other cities in the southeast. Residents of South Fulton acclaimed about the change on social media and thanked the City of South Fulton Police Department and our team for contributing to the communities. The new beat design also received coverage from several news sources, including Fox 5 Atlanta [9].

8 CONCLUSION

In this paper, we presented our work on the City of South Fulton police beat redesign. We propose an optimization framework with the spatial regression model as well as large-scale data analytics. We construct an operational model to predict zone workload using an accurate and tractable linear approximation. The proposed method yields a redesigned zone plan with lower workload variance by only changing eight beats. Currently, we are continuing our partnership with the SFPD. We will continue to observe the police workload in the City of South Fulton as the city and workload grow. If the workload becomes unbalanced once more, we can quickly suggest a new beat design using our already existing methods. As the SFPD continues to grow, they will also hire an information officer that will assist in workload analytics and carry on our workload prediction.




Figure 9: Workload variance with different number of beats.

ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under Grant CMMI-2015787

REFERENCES

- [1] Emile HL Aarts et al. 1987. *Simulated annealing: Theory and applications*. Reidel.
- [2] Deepak Bammi. 1975. Allocation of police beats to patrol units to minimize response time to calls for service. *Computers & Operations Research* 2, 1 (1975), 1–12.
- [3] Dimitris Bertsimas, John Tsitsiklis, et al. 1993. Simulated annealing. *Statist. Sci.* 8, 1 (1993), 10–15.
- [4] Samuel E Bodily. 1978. Police sector design incorporating preferences of interest groups for equality and efficiency. *Management Science* 24, 12 (1978), 1301–1313.
- [5] Jan M Chaiken and Richard C Larson. 1972. Methods for allocating urban emergency units: a survey. *Management Science* 19, 4-part-2 (1972), P110–P130.
- [6] Kenneth Chelst and James P Jarvis. 1979. Estimating the probability distribution of travel times for urban emergency service systems. *Operations Research* 27, 1 (1979), 199–204.
- [7] City of South Fulton. 2017. Destination South Fulton. <https://www.cityofsouthfultonga.gov/2573/Economic-Development>.
- [8] Steven J. D’Amico, Shouu-Jiun Wang, Rajan Batta, and Christopher M. Rump. 2002. A simulated annealing approach to police district design. *Computers & Operations Research* 29, 6 (2002), 667 – 684. [https://doi.org/10.1016/S0305-0548\(01\)00056-9](https://doi.org/10.1016/S0305-0548(01)00056-9)
- [9] Deidra Dukes. 2020. South Fulton making changes to improve office response time. <https://www.fox5atlanta.com/news/south-fulton-making-changes-to-improve-office-response-time>.
- [10] R. S. Garfinkel and G. L. Nemhauser. 1970. Optimal Political Districting by Implicit Enumeration Techniques. *Management Science* 16, 8 (1970), B–495–B–508. <https://doi.org/10.1287/mnsc.16.8.B495>
- [11] Saul I Gass. 1968. On the division of police districts into patrol beats. In *Proceedings of the 1968 23rd ACM National Conference*. ACM, 459–473.
- [12] South Fulton Government. 2020. City of South Fulton GIS. <https://gisdata.fultoncountygga.gov/datasets/city-limits>.
- [13] Linda V Green and Peter J Kolesar. 2004. Anniversary article: Improving emergency responsiveness with management science. *Management Science* 50, 8 (2004), 1001–1014.
- [14] B. Grofman. 1985. Criteria for redistricting: A social science perspective. *UCLA Law Review* 33 (1985), 77–184.
- [15] Lev Grossman, Mark Thompson, Jeffrey Kluger, Alice Park, Bryan Walsh, Claire Suddath, Eric Dodds, Kayla Webley, Nate Rawlings, Feifei Sun, Cleo Brock-Abraham, and Nick Carbone. 2011. The 50 Best Inventions. *Time* (2011). <http://content.time.com/time/subscriber/article/0,33009,2099708-4,00.html>
- [16] Odie Donald II. 2019. City of South Fulton Strategic Plan. <https://www.cityofsouthfultonga.gov/DocumentCenter/View/1608/Draft-City-of-South-Fulton-Strategic-Plan>.
- [17] Richard C Larson. 1972. *Urban police patrol analysis*. Vol. 28. MIT Press Cambridge, MA.
- [18] Richard C. Larson. 1974. A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers & Operations Research* 1, 1 (1974), 67 – 95. [https://doi.org/10.1016/0305-0548\(74\)90076-8](https://doi.org/10.1016/0305-0548(74)90076-8)
- [19] Richard C Larson. 1974. Illustrative police sector redesign in district 4 in Boston. *Urban Analysis* 2, 1 (1974), 51–91.
- [20] Richard C. Larson and Amedeo R. Odoni. 1981. *Urban Operations Research*. Prentice Hall, NJ.
- [21] Andrea Lecchini-Visintini, John Lygeros, and Jan Maciejowski. 2008. Simulated annealing: Rigorous finite-time guarantees for optimization on continuous domains. In *Advances in Neural Information Processing Systems*. 865–872.
- [22] Anuj Mehrotra, Ellis L Johnson, and George L Nemhauser. 1998. An optimization based heuristic for political districting. *Management Science* 44, 8 (1998), 1100–1114.
- [23] G. Mills. 1967. The determination of local government electoral boundaries. *Operations Research Quarterly* 18 (1967), 243–255.
- [24] R. L. Morrill. 1973. Ideal and reality in reapportionment. *Annals of the Association of American Geographers* 63 (1973), 463–477.
- [25] R. L. Morrill. 1976. Redistricting revisited. *Annals of the Association of American Geographers* 66 (1976), 548–556.
- [26] S. S Nagel. 1972. Computers and the law and politics of redistricting. *Polity* 5 (1972), 77–93.
- [27] et al. Niemi, R. G. 1990. Measuring Compactness and the Role of a Compactness Standard in a Test for Partisan and Racial Gerrymandering. *The Journal of Politics* 52(4) (1990), 1155–1181.
- [28] Roger Z Ríos-Mercado. 2020. *Optimal Districting and Territory Design*. Vol. 284. Springer.
- [29] Brian D Ripley. 2005. *Spatial statistics*. Vol. 575. John Wiley & Sons.
- [30] Sherwin Rosen. 1974. Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy* 82, 1 (1974), 34–55. <http://www.jstor.org/stable/1830899>
- [31] Takeshi Shirabe. 2009. Districting Modeling with Exact Contiguity Constraints. *Environment and Planning B: Planning and Design* 36 (11 2009), 1053–1066. <https://doi.org/10.1068/b34104>
- [32] Takeshi Shirabe. 2009. Districting Modeling with Exact Contiguity Constraints. *Environment and Planning B: Planning and Design* 36, 6 (2009), 1053–1066. <https://doi.org/10.1068/b34104> arXiv:<https://doi.org/10.1068/b34104>
- [33] United States Census Bureau. 2018. 2018 Population Estimates Program. <https://factfinder.census.gov>.
- [34] Peter JM Van Laarhoven and Emile HL Aarts. 1987. *Simulated annealing. In Simulated annealing: Theory and applications*. Springer, 7–15.
- [35] W. Vickrey. 1961. On the prevention of gerrymandering. *Political Science Quarterly* 76 (1961), 105– 110.
- [36] Feng-Sheng Wang and Li-Hsunan Chen. 2013. *Heuristic Optimization*. Springer New York, New York, NY, 885–885. https://doi.org/10.1007/978-1-4419-9863-7_411
- [37] J. B. Weaver and S. W. Hess. 1963. Districting Modeling with Exact Contiguity Constraints. *The Yale Law Journal* 72 (1963), 288–308.
- [38] H. P. Yong. 1988. Measuring the compactness of legislative districts. *Legislative Studies Quarterly* 13 (1988), 105–115.