

Contextual Stochastic Optimization for School Desegregation Policymaking

Hongzhao Guan¹, Nabeel Gillani^{2,3}, Tyler Simko⁴, Jasmine Mangat², Pascal Van Hentenryck¹

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

²College of Arts, Media and Design, Northeastern University

³D’Amore-McKim School of Business, Northeastern University

⁴Department of Government, Harvard University

hguan7@gatech.edu, n.gillani@northeastern.edu, tsimko@g.harvard.edu, j.mangat@northeastern.edu, pascal.vanhentenryck@isye.gatech.edu

Abstract

Most US school districts draw geographic “attendance zones” to assign children to schools based on their home address, a process that can replicate existing neighborhood racial/ethnic and socioeconomic status (SES) segregation in schools. Redrawing boundaries can reduce segregation, but estimating expected rezoning impacts is often challenging because families can opt-out of their assigned schools. This paper seeks to alleviate this societal problem by developing a joint redistricting and choice modeling framework, called *redistricting with choices (RWC)*. The RWC framework is applied to a large US public school district to estimate how redrawing elementary school boundaries might realistically impact levels of socioeconomic segregation. The main methodological contribution of RWC is a contextual stochastic optimization model that aims to minimize district-wide segregation by integrating rezoning constraints with a machine learning-based school choice model. The study finds that RWC yields boundary changes that might reduce segregation by a substantial amount (23%) – but doing so might require the re-assignment of a large number of students, likely to mitigate re-segregation that choice patterns could exacerbate. The results also reveal that predicting school choice is a challenging machine learning problem. Overall, this study offers a novel practical framework that both academics and policymakers might use to foster more diverse and integrated schools.

Code — <https://github.com/Plural-Connections/AAAI-2025-Redistricting-with-Choices>

Paper with Supplementary Materials — <https://arxiv.org/abs/2408.12572>

1 Introduction

Schools in the United States remain highly segregated by race and Socio-Economic Status (SES), despite evidence that integration reduces achievement gaps (Reardon, Kalogrides, and Shores 2018; Reardon et al. 2019; Billings, Deming, and Rockoff 2013; Johnson 2019, 2011; Wells and Crain 1994). School segregation is largely driven by high levels of residential segregation across race and SES, as the vast majority of public school students attend the school they are geographically assigned to attend via the “attendance zones” set by their local school district (Monarrez 2020).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The practical difficulty of redrawing attendance boundaries has inspired academic research on automated redrawing methods targeting a number of goals, including integration (e.g. Gillani et al. 2023). However, changing student assignment policies to foster diverse schools often induces changes in how families subsequently pick schools for their children—sometimes yielding “White flight”, or decisions among White and affluent families to opt-out of their newly assigned schools (Reber 2005; Bjerre-Nielsen and Gandil 2020; Macartney and Singleton 2018). *Better anticipating such responses prior to changing boundaries poses a fundamental challenge to designing policies that achieve integration goals in practice*. In other words, the redrawing task requires an interplay between two sub-tasks: choices are first determined conditional on a possible boundary change (“redistricting”), and then a redistricting is produced conditional on those choices. Researchers have previously developed school redistricting algorithms and subsequently estimated how families might select schools in response to changes (Allman et al. 2022)—but their model did not factor in school selections to inform boundary changes themselves.

This paper aims to address this challenge: it develops a novel joint redistricting and choice modeling framework, called *redistricting with choices (RWC)*. The RWC framework proposes a *contextual stochastic constraint program* that minimizes district-wide dissimilarity, and integrates the rezoning constraints and a school choice model. RWC derives the school choice model using a machine learning model, leveraging features of students, Census Blocks, and the district schools.

RWC was evaluated on real data from a large public US school district — the Winston-Salem/Forsyth County Schools in North Carolina, which serves over 50,000 students (22,000 of which are enrolled in elementary schools). The study reveals that the RWC yields an expected 23% decrease in segregation while rezoning more Census Blocks and students than baseline choice models. This additional rezoning is likely necessary to offset a predicted tendency for more students to opt-out of their zoned schools (and perhaps increase segregation as a result). The study also shows that predicting school choice is a challenging multi-class classification problem (F1 score: 0.73). Fortunately, 87% of the time, the classifier identifies a student’s ground-truth school among its top three most probable predictions.

The main contributions of the paper can be summarized as follows.

1. To the authors' knowledge, RWC is the first joint school redistricting optimization model where school choice predictions inform integration-promoting boundary changes.
2. RWC leverages contextual stochastic optimization, and its approximation through the Sample Average Approximation method, to obtain a practical tool to inform the design of school boundaries that foster integration.
3. Experimental results on a large school district show that RWC can substantially reduce segregation, but typically re-assigns more students than other methods.

Together, these contributions demonstrate the important role AI can play in supporting decision-makers seeking to advance equitable access to quality education.

2 Related Work

Much prior work from AI and other fields has demonstrated how attendance zones might be redrawn to reduce racial segregation while respecting local constraints like minimizing travel time to schools (e.g. Smilowitz and Keppler 2020; Heckman and Taylor 1969; Liggett 1973; Clarke and Surkis 1968; Holloway et al. 1975; Diamond and Wright 1987; Gillani et al. 2023). Such tools have also been adopted by academics working in partnership with districts (e.g. Allman et al. 2022). Still, implementing them in practice requires careful thought and attention to the goals of policy-makers and community members. In general, redistricting is a socio-politically fraught topic, particularly when issues of diversity are at stake, because parents fear that boundary changes might impact friend groups (Bridges 2016; Staff 2019), home valuations (Kane, Staiger, and Riegg 2005; Bridges 2016; Black 1999), school quality (Zhang 2008), and many other factors. Beyond this paper, the researcher-practitioner partnership that this work is a part of aims to attend to these concerns while also fostering integration.

Researchers have also previously sought to predict school choice. Anticipating school choice is important for planning teams in school districts. Prior work on predicting school choice has developed (mostly linear) demand choice models in centralized school assignment settings (e.g Pathak and Shi 2017; Shi 2022). However, these studies have not applied such models in the context of simultaneously adjusting geographic assignment policies (like attendance boundaries). A notable exception is (Allman et al. 2022), yet as noted earlier, this study is still preliminary given it estimates school choices post-redistricting instead of factoring them into boundary changes. This paper aims at advancing this work by estimating family choices and using those estimates to shape the boundary change process.

Note that researchers in other settings have explored the integration of choice models in optimization problems (e.g., optimizing transportation network and routes, which requires designing transit network conditional on riders' choices, and vice versa (Basciftci and Van Hentenryck 2022; Guan, Basciftci, and Van Hentenryck 2024)). Such models often computationally challenging in practice.

3 Research Setting



Figure 1: Maps of WS/FCS. Left shows race/ethnicity (computed via (Kenny 2022))—lighter yellow means higher proportion Black/Hispanic); right shows the distribution of the continuous, geographically-defined Socio-Economic Status (SES) index measure values (described more in Section 4—darker purple means lower SES).

The research partner for this study is Winston-Salem/Forsyth County Schools (WS/FCS), a public school district serving students from pre-kindergarten to twelfth grade in North Carolina. With more than 50,000 students, WS/FCS is the 4th largest school district in North Carolina and the 81st largest in the United States. WS/FCS students are racially diverse, with approximately 33% White students, 30% Black students, 30% Hispanic students, 6% reporting multiple races, and 3% Asian students. The district operates 81 schools (42 elementary). Despite district-wide diversity, residential and school enrollment patterns are still highly segregated. Figure 1 visualizes this segregation, highlighting patterns that make WS/FCS one of the most segregated districts in the state. As a recent recipient of a US Department of Education's "Fostering Diverse Schools" grant (Jacobson 2023), there is renewed interest in WS/FCS to explore how boundary changes might enhance SES diversity in schools. These would be the first significant changes to district boundaries since the end of court-ordered desegregation in the 1990s (Garms and Starke 2023). This paper names the district collaborator (with their permission) to ground the empirical explorations and findings in relevant historical and present-day context and inform more thoughtful applications of AI for social impact.

Most WS/FCS students attend their zoned school, and therefore, the focus of the Fostering Diverse Schools work is largely on redrawing boundaries based on underlying patterns of residential segregation, without trying to anticipate school choice. Still, the district has an active choice program, meaning exploring how choice patterns might alter the effects of boundary changes is essential for the sound implementation of new policies. WS/FCS families are allowed to choose either the school they are residentially (default) assigned to; a magnet school; any school inside a pre-set, contiguous group of schools called a "choice zone" (with guaranteed transportation); or any school outside of their choice zone (in which case, they must provide their own transportation). The district does not use deferred acceptance to

centralize student-school matching. Typically, students are able to select any school they wish to attend (and in cases of schools being oversubscribed, the district is usually still able to accommodate them). Therefore, this paper does not simulate deferred acceptance, top-trading cycles, or another typical centralized student assignment mechanism; students are simply assigned to schools by sampling their most likely choice from a probability distribution over schools.

4 Problem and Methodology

This section presents the problem setting and the methodological approach of the paper. The overall goal is to choose a zoning for a public school district that minimizes school dissimilarities with respect to a group of students. The main methodology contribution is to include a choice model into the optimization, i.e., capturing how students will respond to changes in their allocated schools. From a technical standpoint, the methodology uses contextual stochastic optimization (Sadana et al. 2024) as a framework and constraint programming for modeling and solving the resulting optimization problem. Figure 2 summarizes the overall methodology, and Table 1 summarizes the notation used in the paper.

Data and Notations

The redistricting problem seeks to assign Census Blocks to schools to minimize some metrics. Such an assignment is called a *zoning*. A census block is the smallest geographic area determined and publicly released by the United States Census Bureau. Each census block typically contains many households, but it can have no population and only include geographical features such as rivers and mountains.

Four years (2019—2023) of anonymized, geocoded student data were obtained from Winston-Salem/Forsyth County Schools through a signed data-sharing agreement, with University IRB approval. During the 2022-2023 school year, the district operated 41 public elementary schools across 6,373 Census Blocks, and enrolled 22,302 students.

The set of Census Blocks is denoted by \mathcal{B} and the set of schools by \mathcal{SH} . In the following, a bar notation, e.g., $\bar{t}_{b,s}$, represents the value of a quantity in the dataset. \bar{s}_b denotes the school assigned of each census block b in the dataset zoning. The census block of student n is denoted by \bar{b}_n . The school assigned to student n , denoted by \bar{s}_n^{zone} , can be derived from \bar{b}_n and \bar{s}_b . However, students are allowed to opt-out to any of 40 other public schools in the district, and \bar{s}_n^{actual} represents the school in which the student is actually enrolled and \bar{c}_s represents the actual student enrollment at each school s . Each student n is assigned a Socio-Economic Status (SES) index, denoted by \bar{e}_n , based on their residential area. This index is categorized into three levels: lower (0), medium (1), and higher (2). Its definition is presented in more detail in (Gillani and Simko 2024) and is similar to one used by peer districts (Quick 2016; Hawkins 2018). At a high level, the measure is created by first producing a geographically-defined, continuous index measure comprised of a number of American Community Survey variables (median household income, adult educational attainment levels, etc.), and then thresholding this continuous

Notation	Definition
Sets:	
$\mathcal{SH}, \mathcal{SH}_{mgt}$	Schools and magnet schools considered in this study ($\mathcal{SH}_{mgt} \subseteq \mathcal{SH}$, $ \mathcal{SH} = 41$, $ \mathcal{SH}_{mgt} = 8$).
\mathcal{SH}_r	The nearest r schools to a student's residential block.
\mathcal{B}	Census Blocks considered in this study ($ \mathcal{B} = 6373$).
\mathcal{Z}	The set of all feasible assignments.
Vectors:	
$\mathbf{f}_n, \mathbf{d}_n$	The static and dynamic feature vectors of student n .
\mathbf{x}_n	The contextual information and feature vector of student n , concatenated by \mathbf{f}_n and \mathbf{d}_n .
\mathbf{y}	The ground-truth labels, indicating the schools students actually attend in the dataset, i.e., $\mathbf{y}_n = \bar{s}_n^{actual}$.
\mathbf{z}	A zoning, constructed by decision variables $z_{b,s}$.
\mathbf{S}	A vector of random variables S_n .
Decision Variables:	
$z_{b,s}$	Binary variable indicates if a census block b is zoned to school s .
d, d^i	Continuous variable indicates the district-wide dissimilarity, dynamically based on \mathbf{z} . i indicates scenarios for SAA.
c_s, c_s^i	Integer variable indicates the number of students in school s , dynamically based on \mathbf{z} . i indicates scenarios for SAA.
g_s, g_s^i	Integer variable indicates the number of lower-SES students in school s , dynamically based on \mathbf{z} . i indicates scenarios for SAA.
Parameters:	
I	The number of SAA (Sample Average Approximation) scenarios.
τ	The maximum increment ratio on travel time, $\tau \in [0, 1]$.
α	The maximum change ratio in school population, $\alpha \in [0, 1]$.
Constants:	
N	The number of student in the study ($N = 22, 302$).
\bar{g}_{total}	The number of student in the target group. In this paper, the target group is students in lower-SES category ($\bar{g}_{total} = 9, 039$).
\bar{e}_n	The SES category of student n , determined by their residential area. $\bar{e}_n = 0$ means student n belongs to the lower-SES category.
$\bar{t}_{b,s}$	The estimated driving time from census block b to school s .
\bar{b}_n	The residential census block of student n .
\bar{c}_s	The number of current students in school s , based on realistic data.
\bar{s}_b	The current zoned school to block b , based on realistic data.
$\bar{s}_n^{zone}, \bar{s}_n^{actual}$	The ground-truth zoned (residentially-assigned school) and actual schools for student n , respectively, based on realistic data.
Others:	
S_n	Random variable indicating the choice of school of student n .
s_n^{zone}	A variable that is used to indicate the zoned school for student n , useful for constructing \mathbf{x}_n .
$A_n^i(\mathbf{z})$	The school chosen by student n for zoning \mathbf{z} . i indicates scenarios for SAA.
C, C^f, C^{rb}, C^{ml}	Choice models employed in this study, "f", "rb", and "ml", stand for follow, rule-based, and machine learning, respectively.
C^{ml-l}, C^{ml-x}	Continued from the above, machine learning based choice models, "l" and "x" stand for multinomial-logit and XGboost, respectively.

Table 1: Notation used throughout this study.

measure to create the discrete categories described above: students living in Census Blocks in the bottom, middle, and top third of the distribution for this measure are considered lower, medium, or higher-SES, respectively.

Decision Variables and Objective Function

The decision variables capture the proposed zoning, i.e., binary variable $z_{b,s}$ represents whether block b is assigned to school s are represented by binary decision variables. This collection of decision variables is collected in a zoning vector \mathbf{z} . The auxiliary variables c_s and g_s represent the number of students and the number of students with lower-SES scores assigned to school s . Zonings must satisfy several important constraints that are described later in the paper, and the set of feasible zonings is denoted by \mathcal{Z} .

The optimization model aims at finding a zoning that minimizes the district-wide dissimilarity metric (Massey and Denton 1988), i.e.,

$$d = \frac{1}{2} \sum_{s \in S} \left| \frac{g_s}{\bar{g}_{total}} - \frac{c_s - g_s}{N - \bar{g}_{total}} \right| \quad (1)$$

where constant \bar{g}_{total} represents the total number of students with lower-SES scores across all schools (and so, the

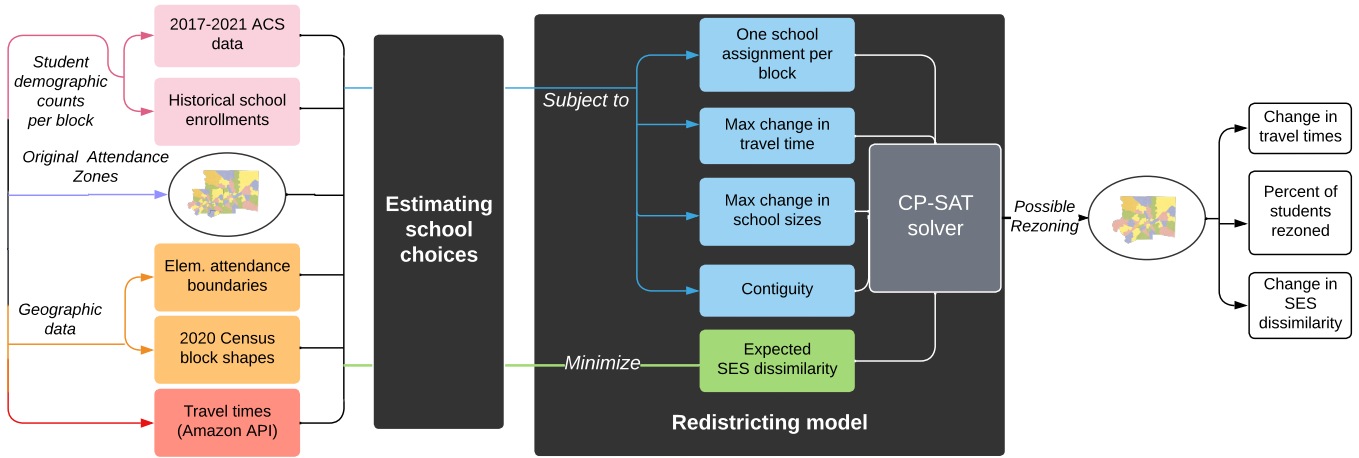


Figure 2: Diagram for the RWC framework, adapted from (Gillani et al. 2023).

objective is to minimize the segregation of lower-SES students from their medium and higher-SES counterparts). The dissimilarity index of a school s ranges from 0 (perfect integration—the demographic mix of each school reflects district-wide proportions) to 1 (total segregation—all lower-SES students in the district attend a single school).

Contextual Stochastic Optimization (CSO)

One of the key challenges of the redistricting problem is to capture student choices for schools. For instance, a student might attend school 1 when assigned school 1, but opt-out to school 3 when assigned to school 2. These decisions, of course, are not available directly; to overcome this limitation, this paper explores the use of machine learning and contextual stochastic optimization.

Let S_n be a random variable representing the school student n attends for a proposed zoning and \mathbf{S} be the vector of these random variables. For a zoning \mathbf{z} , a realization of \mathbf{S} defines the auxiliary variables c_s and g_s defined earlier and hence the district-wide dissimilarity metric. This is captured by the function $dis-score(\mathbf{s})$ which computes the district-wide dissimilarity score from a realization of \mathbf{S} .

The redistricting optimization problem can then be formalized as the following stochastic optimization problem:

$$\min_{\mathbf{z} \in \mathcal{Z}} \mathbf{E}_{\Pr[\mathbf{S}|\mathbf{z}]} [dis-score(\mathbf{S})] \quad (2)$$

It assumes that students independently choose schools:

$$\Pr[\mathbf{S} | \mathbf{z}] = \prod_{n=1, \dots, N} \Pr[S_n | \mathbf{z}] \quad (3)$$

Again, the distribution $\Pr[S_n | \mathbf{z}]$ is not readily available. In reality, students may influence one another, so future work should explore lifting this independence assumption.

To approximate (2), CSO combines contextual information about students and schools. In particular, the contextual information of student n , denoted by \mathbf{x}_n , comprises personal information, data from their residential census block, as well as the current school-related information (e.g., composition

of the student body). The distribution $\Pr[S_n | \mathbf{z}]$ is then approximated by $\Pr[S_n | \mathbf{z}, \mathbf{x}_n]$ and the CSO problem becomes:

$$\min_{\mathbf{z} \in \mathcal{Z}} \mathbf{E}_{\prod_n \Pr[S_n | \mathbf{z}, \mathbf{x}_n]} [dis-score(\mathbf{S})] \quad (4)$$

Note that, in this formulation, school choices are solely based on information available before actual school attendance. In other words, \mathbf{x}_n is only created using historical data from the 2019-2020 through 2021-2022 school years and the status-quo zoning information for the 2022-2023 school year (\bar{s}_n^{zone}). Those features capture information that previous studies, such as (Pathak and Shi 2017), have shown to provide valuable insights into students' and families' school choices. This includes details such as the programs offered by schools, school ratings, travel times to schools, and other relevant factors.

The CSO problem (2) is approximated with the Sample Average Approximation (SAA) method (Kleywegt, Shapiro, and Homem-de Mello 2002). The solution method generates I scenarios, each of which defines a choice function for each student. More precisely, for each student n and each scenario i , $A_n^i(\mathbf{z})$ denotes the school chosen by student n for zoning \mathbf{z} . The optimization becomes:

$$\min_{\mathbf{z} \in \mathcal{Z}} \sum_{i \in [I]} [dis-score((A_1^i(\mathbf{z}), \dots, A_N^i(\mathbf{z})))] \quad (5)$$

Since each student chooses their school independently, this paper also uses $A_n^i(s)$ to denote the school chosen by student n when assigned to school s .

The Constraint Programming Model

Figure 3 depicts a constraint programming model for formulation (5), where d^i , c_s^i , and g_s^i represent the auxiliary variables computed for a scenario i given a proposed zoning. The model features the following constraints:

1. **School Sizes and Population Increases.** Equations (6c) and (6d) indicate how the total (c_s^i) and lower-SES populations (g_s^i) of schools are derived based on the value of

$$\min \sum_{i=1}^I d^i \quad (6a)$$

$$\text{s.t. } d^i = \frac{1}{2} \sum_{s \in \mathcal{S}} \left| \frac{g_s^i}{\bar{g}_{total}} - \frac{c_s^i - g_s^i}{N - \bar{g}_{total}} \right| \quad (6b)$$

$$i = 1, \dots, I$$

$$c_s^i = \sum_{n=1}^N \sum_{s' \in \mathcal{S}} z_{\bar{b}_n, s'} \cdot \mathbb{1}[A_n^i(s') = s] \quad (6c)$$

$$\forall s \in \mathcal{S}, i = 1, \dots, I$$

$$g_s^i = \sum_{n=1}^N \mathbb{1}[\bar{e}_n = 0] \sum_{s' \in \mathcal{S}} z_{\bar{b}_n, s'} \cdot \mathbb{1}[A_n^i(s') = s] \quad (6d)$$

$$\forall s \in \mathcal{S}, i = 1, \dots, I$$

$$\bar{c}_s^i \cdot (1 + \alpha) \geq c_s^i \geq (1 - \alpha) \cdot \bar{c}_s^i \quad (6e)$$

$$\forall s \in \mathcal{S}, i = 1, \dots, I$$

$$1 = \sum_{s \in \mathcal{S}} z_{b, s} \quad \forall b \in \mathcal{B} \quad (6f)$$

$$\bar{t}_{b, \bar{s}_b} \cdot (1 + \tau) \geq \sum_{s \in \mathcal{S}} z_{b, s} \cdot \bar{t}_{b, s} \quad \forall b \in \mathcal{B} \quad (6g)$$

$$\text{ENSURE CONTIGUITY}(b) \quad \forall b \in \mathcal{B} \quad (6h)$$

$$a_{b, s} \in \{0, 1\} \quad \forall b \in \mathcal{B}, s \in \mathcal{S} \quad (6i)$$

Figure 3: The Constraint Programming Model.

$z_{b, s}$. Constraint (6e) makes sure that the school populations do not exceed, or fall below, some thresholds specified by a parameter α (set to 0.15 based on the same parent survey as described above). Note that constraint (6e) needs to be satisfied for all scenarios $i = 1, \dots, I$.

2. **Travel Time Increases.** Constraint (6g) ensures that driving times to assigned schools do not exceed some relative threshold τ (set to 0.5, inferred from a prior parent survey administered as a part of (Gillani et al. 2023)). Driving times $\bar{t}_{b, s}$ were estimated from the centroid of each Census block to each elementary school.
3. **Contiguity.** Constraint (6h) ensures reasonable zoning shapes, where blocks that are contiguous with respect to their assigned schools remain contiguous in any new zoning (see the Supplementary Materials of (Gillani et al. 2023) for additional details on how this is enforced, inspired by (Mehrotra, Johnson, and Nemhauser 1998)).

Assuming that students in their present schools (\bar{s}_n^{actual}) do not want another school, the model always has a feasible solution, since it suffices to assign students to their school.

Constraint programming was used for its modeling flexibility, though other approaches are possible (Gurnee and Shmoyes 2021; McCartan and Imai 2023; Autrey et al. 2019).

Machine Learning for School Choice Modeling

The contextual distributions $\Pr[S_n | \mathbf{z}, \mathbf{x}_n]$ are approximated by supervised machine learning (e.g., boosted trees or multinomial logit models). The learning task can be seen as a multi-class classification problem, where each class represents a particular school that a student might attend. The model is trained from a dataset $\{(\mathbf{f}_k, \mathbf{d}_k), \mathbf{y}_k\}_k$ where, for each training instance k , \mathbf{f}_k are static features related to demographics and geography, \mathbf{d}_k are dynamic features related to school zoning (together they form the feature matrix), and \mathbf{y}_k is the students' current school choices (i.e., \bar{s}_n^{actual} for student n in the 22-23 school year). The Supplementary Materials include additional technical details on these features.

5 Computational Results

This section presents the computational results for RWC. It first focuses on evaluating the predictive capabilities of \mathcal{C}^{ml} for student school choices. The experiments employ two widely-adopted machine learning algorithms: multinomial logit and XGBoost (Chen and Guestrin 2016). For comparative analysis, two baseline choice models are introduced. The performance of these models is then assessed in the multi-class classification task introduced earlier. Secondly, the optimization results are presented, showing the benefit of using RWC on school redistricting. Finally, this section analyzes the performance stability of the SAA method.

The results are produced using Python 3.10, the scikit-learn (1.5.1) and XGBoost (2.1.1) packages for machine learning, and the “CP-SAT” solver for constraint programming, with 8 threads for each run. All computations were performed using the first CPUs available on a Linux High Performance Computing cluster with Intel Xeon Gold, Platinum, and AMD Epyc CPUs.

Baselines for Choice Modeling

This section specifies two baselines for choice modeling. The first baseline \mathcal{C}^f simply assigns students to their zoned school. Since approximately 2/3 of elementary students already follow this policy in the dataset, it is a coarse simplification to apply this to all students. It is also a valuable baseline for the machine learning models, since they might trivially “learn” to predict that students will select their zoned schools and thus be correct approx. 2/3 of the time.

The second baseline \mathcal{C}^{rb} uses frequency information in the dataset, i.e., how frequently students attend or opt-out of their zoned schools. As mentioned before, students attend their zoned school about 65% of the time; a nearby magnet school 20% of the time; and another nearby school the rest of the time, since families tend to prefer lower travel times to school (Gillani et al. 2023). The choice model is given by

$$\Pr[S_n = s | \mathbf{x}_n] = \begin{cases} 0.65 & \text{if } s = s_n^{zone} \\ \frac{0.2}{|\mathcal{S}\mathcal{H}_{mgt} \cap \mathcal{S}\mathcal{H}_{near,12}|} & \text{if } s \in \mathcal{S}\mathcal{H}_{mgt} \cap \mathcal{S}\mathcal{H}_{near,12} \\ 0.03 & \text{if } s \in \mathcal{S}\mathcal{H}_{near,5} \end{cases} \quad (7)$$

where $\mathcal{S}\mathcal{H}_{mgt}$ represents the magnet schools, and $\mathcal{S}\mathcal{H}_{near,r}$ denotes the r closest schools for student n .¹ Note that s_n^{zone} is a variable, which is critical to construct \mathbf{x}_n .

¹The number 12 was determined as the minimum number of

Choice Model	acc.	top 3 acc.	top 5. acc	Macro Averaged			Weighted Averaged		
				precision	recall	f1-score	precision	recall	f1-score
C^f (baseline 1)	0.6533	-	-	0.6534	0.6522	0.6428	0.6651	0.6533	0.6498
C^{rb} (baseline 2)	0.4606	0.7013	0.7686	0.5124	0.4610	0.4710	0.5338	0.4606	0.4810
C^{ml-l} (logit)	0.6917	0.8516	0.9060	0.6859	0.6854	0.6827	0.6925	0.6917	0.6894
C^{ml-x} (XGboost)	0.7421	0.8748	0.9200	0.7360	0.7354	0.7345	0.7422	0.7421	0.7410

Table 2: Evaluation of different choice models. For C^{ml-l} and C^{ml-x} , their top 3 acc. and top 5 acc. are averaged across the 10-fold validation test sets. All other metrics are computed on the summed confusion matrices across all cross-validation fold.

Choice Modeling Results

This section reports the machine learning results and their comparison with the baselines C^f and C^{rb} . The machine learning models are denoted by C^{ml-l} and C^{ml-x} for the multinomial logit and XGboost, respectively. Both models are evaluated using 10-fold cross-validation, randomly holding out 10% of students each time as the test set. For C^{ml-l} , scikit-learn defaults are used for all parameters except for the max iteration amount, which is set to 2,000. C^{ml-x} is trained with the “multi” objective and “mlogloss” metric; the model implements early stopping with a validation set (15% of the training set) to help prevent over-fitting; and the learning rate and the maximum depth of a tree are tuned to 0.1 and 6 after hyper-parameters tuning with [0.08, 0.1, 0.12] and [4, 6, 8, 10, 12] on the validation set, respectively.

The model outputs \mathbf{y}^{pred} are compared to the ground truth \mathbf{y} for test-set students to evaluate accuracy. Table 2 shows accuracy metrics for the different choice models. In general, C^{ml-x} outperforms all the other models. Still, the accuracy and macro-averaged F1 scores of 0.74 and 0.73, respectively, suggest that school choice prediction is a challenging task. One reason for this may be the large classification space of 41 schools. Another is that many unobservable factors might influence a family’s decision to select any particular school. This is confirmed by prior research in a district where families submit ranked preferences for schools: the study found that multinomial logit models were correct less than 50% of the time in predicting a student’s top-choice school (Pathak and Shi 2017). The current paper offers a large improvement over this previous finding – improvements that may be attributed to the quality of the dataset and feature engineering. Details on the specific features included in the models are included in the Supplementary Materials. Exploring other features is also a valuable direction for future research.

Interestingly, the top-3 and top-5 accuracy results are much higher for C^{ml-x} : 87% of the time, the student’s actual school selection is captured by the top-3 highest-probability schools indicated by the model. For top-5, this value is at 92%. Inspecting the test set probabilities assigned to the top-3 schools across the 10 cross-validation folds shows that, on average, the highest-probability school receives a probability mass of 76%, followed by 10% for the second highest-probability school, followed by 4% for the third (with standard errors close to zero).

nearby schools needed to be included to ensure, for each student, that at least one of them is a magnet school.

Redistricting with choice

The CSO of RWC was evaluated for the following three choice models: C^f , C^{rb} , and C^{ml-x} . These are denoted by R for Redistricting), FR (for Frequency-based Redistricting), and RWC respectively. C^{ml-l} was not considered due to its lower performance in comparison C^{ml-x} . Before the redistricting steps, C^{ml-x} is retrained using the same parameters as the best-fitting model described earlier, but with an 85%-15% split for the training and validation sets (no test set is required since the model’s primary use-case now is inference). For the SAA method, the number of scenarios I is set to 30 and the CP solver is configured to terminate after a 12-hour solving period (none were able to prove optimality).

Table 3 shows the redistricting results from these models. Note that none of the CP models were solved to optimality within 12 hours due to the complexity of the problem. The table also includes statistics for the current (status quo) assignment. Model R represents the “best-possible” redistricting by ignoring that students may opt-out from their boundary-assigned schools. It decreases dissimilarity by 32% (from 0.596 to 0.403) across the district. RWC, which integrates the machine learning choice model and is grounded in a much more realistic setting, produces a comparably impressive reduction in SES segregation of 23%. This reduction is nearly 2x larger than reductions discovered in previous school redistricting work focused on fostering racial/ethnic integration, but without choice modeling (Gillani et al. 2023). Interestingly, RWC also outperforms FR, which uses a much less realistic choice model. Travel times across all models, as well as the status quo, are similar—suggesting that students opting out of their zoned schools still generally select ones nearby.

Further inspecting the results in Table 3 suggests that the models appear to be making different decisions about which students would opt out of their zoned schools. As expected, by virtue of how the rule-based choice model is defined for FR, approximately 1/3 of students are estimated to opt-out of their zoned school even after redistricting (which closely mirrors ground-truth opt-out rates). In the RWC model, more than 40% are estimated to opt out—however, as shown in Supplementary Materials, opt-out rates under this model seem to be consistent across demographics. Interestingly, as the columns labeled “# Rezoned” show, RWC also tends to rezone more Census Blocks, lower-SES students, and students overall compared to the other models. These trends may reflect the need to make more drastic boundary changes to account for more realistic patterns of choice that might

Method	# SAA Scenarios (I)	Dissimilarity (standard error)	# Rezoned (% Rezoned)				Average # Students opting out of zoned school	Average driving time (min)
			Lower-SES students	All students	Census Blocks			
Current	-	0.596 (-)	-	-	-	7,731 (34.67%)	7.54	
R	-	0.403 (-)	1,870 (20.69%)	5,369 (24.07%)	1,902 (29.84%)	5,731 (25.70%)	7.44	
FR	30	0.462 (4.0×10^{-4})	1,870 (20.69%)	4,397 (19.72%)	1,512 (23.73%)	7,433 (33.33%)	7.53	
RWC	30	0.459 (4.6×10^{-4})	3,465 (38.33%)	7,864 (35.26%)	2,307 (36.20%)	9,227 (41.37%)	7.88	

Table 3: Results from different redistricting methods. For FR and RWC, dissimilarity values are averaged over scenarios (with resulting dissimilarity index standard errors in parenthesis). Note that R is deterministic. Interestingly, both R and FR resulted in the rezoning of 1,870 lower-SES students. This similarity is coincidental (the maps in Supplementary Materials are different).

re-segregate schools, as choice sometimes does (Whitehurst 2017; Candipan 2019). Maps with results from the different redistricting models can be found in the Supplementary Materials. The findings show that RWC is not trivially replicating R (predicting students will attend their zoned school).

The stability of the SAA method for RWC is evaluated through additional experiments. Specifically, 100 RWC runs are conducted with $I = 30$, resulting in an average outcome of 0.466 dissimilarity score and a standard error of 0.0012. Further investigation involved increasing I to 50, with 100 RWC runs performed under these conditions. The results showed a higher average dissimilarity of 0.495 and a standard error of 0.0049—i.e., on average, the results are less effective in reducing segregation than when $I = 30$. This may be due to the fact that while 61 out of the 100 runs produce results below 0.47, which aligns with findings when $I = 30$, the remaining runs result in much higher dissimilarity values, closer to the observed ground truth of 0.6. This behavior suggests that the observed variations are likely due to challenges associated with optimizing the CSO model as I increases, given that a larger value of I also translates into a more complex optimization problem. Future research may develop more efficient methods for solving the RWC model.

6 Discussion and Conclusion

This paper presents RWC, a novel joint redistricting and choice modeling framework that uses a *contextual stochastic constraint program* to minimize district-wide SES segregation. RWC derives the school choice model using a machine learning model, leveraging features of students, Census Blocks, and district schools. With a realistic choice model, RWC results in large reductions (23%) in SES segregation, suggesting it may help foster more integrated schools even after accounting for families who alter their school selections in response to such changes. RWC can thus support both researchers and practitioners in realizing the renewed promise of fostering diverse schools (Jacobson 2023).

Still, the paper has a number of limitations, which may serve as important directions for future work. Perhaps the most obvious is the choice model’s performance. Anticipating school choice has been a challenging task for other researchers like (Pathak and Shi 2017). Nevertheless, an exciting opportunity for future work includes identifying more accurate prediction approaches – for example, through syn-

thetic or historical data augmentation; a richer feature space that additionally factors in qualitative preferences (e.g. parents’ online school reviews (Gillani et al. 2021)); and more sophisticated machine learning models.

Another limitation is that this study only models within-district choice. Families may leave district schools for private, public charter, or other options following a student assignment policy change (Macartney and Singleton 2018; Reber 2005; Bjerre-Nielsen and Gandil 2020; Mervosh 2021). This out-of-district choice is more difficult to model due to limited ground truth data. Still, finding ways to anticipate this choice may yield a more accurate picture of a given redistricting’s expected impact on integration.

Even the present within-district redistricting-with-choice approach might be improved by iterating the choice and redistricting models. At present, this study does not simulate a true “game” between boundary changes / impacts on school populations and the choices that families might make. The study estimates choices as a function of changes to a student’s zoned school assignment; but in practice, as prior work on neighborhood “tipping” (Card, Mas, and Rothstein 2008) has suggested, a family’s choice may also depend on how aspects of a school that can only be measured post-rezoning (like their demographics) might change. Modeling this more complex interplay would require more computationally demanding bi-level optimization approaches.

Finally, future work may explore how the RWC model makes decisions, and why it tends to produce more drastic boundary changes. While the model achieves larger reductions in segregation, it leads to more students being rezoned and more opt-outs. This reflects prior literature like (Bjerre-Nielsen and Gandil 2020) suggesting that families may opt-out of new school assignments if they are unfavorable. Understanding the predictors of opt-outs through interpretable machine learning methods like (Lundberg et al. 2020) may help districts better understand and prepare for community responses to integration policies.

Fostering integrated schools is a perennial challenge across US public school districts with large implications for access to quality educational environments and future life outcomes. Anticipating the potential impacts of integration policies requires both factoring in top-down policy priorities and anticipating bottom-up responses to potential changes. RWC offers one starting point for doing so in order to advance more equitable policies in public education.

Acknowledgements

This work was partially supported by the US Department of Education's Fostering Diverse Schools Grant, the Overdeck Family Foundation, and NSF AI Institute for Advances in Optimization Grant 2112533.

References

- Allman, M.; Ashlagi, I.; Lo, I.; Love, J.; Mentzer, K.; Ruiz-Setz, L.; and O'Connell, H. 2022. Designing School Choice for Diversity in the San Francisco Unified School District. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, 290–291.
- Autrey, E.; Carter, D.; Herschlag, G.; Hunter, Z.; and Mattingly, J. C. 2019. Metropolized forest recombination for monte carlo sampling of graph partitions. *arXiv preprint arXiv:1911.01503*.
- Basciftci, B.; and Van Hentenryck, P. 2022. Capturing Travel Mode Adoption in Designing On-Demand Multimodal Transit Systems. *Transportation Science*, 1–25.
- Billings, S. B.; Deming, D. J.; and Rockoff, J. 2013. School Segregation, Educational Attainment, and Crime: Evidence from the End of Busing in Charlotte-Mecklenburg. *The Quarterly Journal of Economics*, 129(1): 435–476.
- Bjerre-Nielsen, A.; and Gandil, M. H. 2020. Attendance boundary policies and the limits to combating school segregation. *Working paper*.
- Black, S. E. 1999. Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics*, 114(2): 577–599.
- Bridges, K. 2016. Eden Prairie Public Schools: Adapting to Demographic Change in the Suburbs. The Century Foundation.
- Candipan, J. 2019. Neighbourhood change and the neighbourhood-school gap. *Urban Studies*, 56(15).
- Card, D.; Mas, A.; and Rothstein, J. 2008. Tipping and the Dynamics of Segregation. *The Quarterly Journal of Economics*, 123(1): 177–218.
- Chen, T.; and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.
- Clarke, S.; and Surkis, J. 1968. An operations research approach to racial desegregation of school systems. *Socio-Economic Planning Sciences*, 1(3): 259–272.
- Diamond, J. T.; and Wright, J. R. 1987. Multiobjective analysis of public school consolidation. *Journal of Urban Planning and Development*, 113(1): 1–18.
- Garms, L.; and Starke, S. 2023. Rooted in Race: Our Community's History of School Integration. The Winston-Salem Foundation.
- Gillani, N.; Beeferman, D.; Vega-Pourheydarian, C.; Overney, C.; Van Hentenryck, P.; and Roy, D. 2023. Redrawing attendance boundaries to promote racial and ethnic diversity in elementary schools. *Educational Researcher*, 52(6): 348–364.
- Gillani, N.; Chu, E.; Beeferman, D.; Eynon, R.; and Roy, D. 2021. Parents Online School Reviews Reflect Several Racial and Socioeconomic Disparities in K-12 Education. *AERA Open*, 7(1): 1–16.
- Gillani, N.; and Simko, T. 2024. School Districts Can Reduce Segregation by Redrawing Attendance Zones, but the Choice to Target Race/Ethnicity or SES Matters. *Working paper*.
- Guan, H.; Basciftci, B.; and Van Hentenryck, P. 2024. Path-Based Formulations for the Design of On-demand Multimodal Transit Systems with Adoption Awareness. *INFORMS Journal on Computing*.
- Gurnee, W.; and Shmoys, D. B. 2021. Fairmandering: A column generation heuristic for fairness-optimized political districting. *arXiv: 2103.11469*.
- Hawkins, B. 2018. 78207: America's Most Radical School Integration Experiment. The 74 Million.
- Heckman, L. B.; and Taylor, H. M. 1969. School rezoning to achieve racial balance: a linear programming approach. *Socio-Economic Planning Sciences*, 3(2): 127–133.
- Holloway, C. A.; Wehrung, D. A.; Zeitlin, M. P.; and Nelson, R. T. 1975. An interactive procedure for the school boundary problem with declining enrollment. *Operations Research*, 23(2): 191–206.
- Jacobson, L. 2023. Feds Award Millions to School Districts to Address 'Tricky' Issue of Integration. The 74 Million.
- Johnson, R. C. 2011. Long-run Impacts of School Desegregation & School Quality on Adult Attainments. *NBER Working Paper No. 16664*.
- Johnson, R. C. 2019. *Children of the dream: Why school integration works*. Basic Books.
- Kane, T. J.; Staiger, D. O.; and Riegg, S. K. 2005. School Quality, Neighborhoods and Housing Prices: The Impacts of school Desegregation. *NBER Working Paper No. 11347*.
- Kenny, C. T. 2022. *censable: Making Census Data More Usable*. <https://christophertkenny.com/censable/>, <https://github.com/christopherkenny/censable>.
- Kleywegt, A. J.; Shapiro, A.; and Homem-de Mello, T. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on optimization*, 12(2): 479–502.
- Liggett, R. S. 1973. The application of an implicit enumeration algorithm to the school desegregation problem. *Management Science*, 20(2): 159–168.
- Lundberg, S. M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J. M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; and Lee, S.-I. 2020. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2: 56–67.
- Macartney, H.; and Singleton, J. D. 2018. School boards and student segregation. *Journal of Public Economics*, 164: 165–182.
- Massey, D. S.; and Denton, N. A. 1988. The Dimensions of Residential Segregation. *Social Forces*, 67(2): 281–315.

McCartan, C.; and Imai, K. 2023. Sequential Monte Carlo for sampling balanced and compact redistricting plans. *The Annals of Applied Statistics*, 17(4): 3300–3323.

Mehrotra, A.; Johnson, E. L.; and Nemhauser, G. L. 1998. An Optimization Based Heuristic for Political Districting. *Management Science*, 44(8): 1021–1166.

Mervosh, S. 2021. In Minneapolis Schools, White Families Are Asked to Help Do the Integrating. *The New York Times*.

Monarrez, T. E. 2020. School attendance boundaries and the segregation of public schools in the us. *American Economic Journal: Applied Economics*.

Pathak, P. A.; and Shi, P. 2017. How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice. *NBER Working Paper No. 24017*.

Quick, K. 2016. Chicago Public Schools: Ensuring Diversity in Selective Enrollment and Magnet Schools. *The Century Foundation*.

Reardon, S. F.; Kalogrides, D.; and Shores, K. 2018. The Geography of Racial/Ethnic Test Score Gaps. *Stanford Center for Education Policy Analysis, Working Paper No. 16-10*.

Reardon, S. F.; Weathers, E.; Fahle, E.; Jang, H.; and Kalogrides, D. 2019. Is separate still unequal? New evidence on school segregation and racial academic achievement gaps. *Stanford Center for Education Policy Analysis*.

Reber, S. J. 2005. Court-Ordered Desegregation: Successes and Failures Integrating American Schools since Brown versus Board of Education. *Journal of Human Resources*, 40(3): 559–590.

Sadana, U.; Chenreddy, A.; Delage, E.; Forel, A.; Frejinger, E.; and Vidal, T. 2024. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*.

Shi, P. 2022. Optimal Priority-Based Allocation Mechanisms. *Management Science*, 68(1).

Smilowitz, K.; and Keppler, S. 2020. On the use of operations research and management in public education systems. *Pushing the boundaries: Frontiers in impactful OR/OM research*, 84–105.

Staff, B. S. 2019. In Howard County, a ‘courageous’ plan to redraw school boundaries tests community’s commitment to diversity. *The Baltimore Sun*.

Wells, A. S.; and Crain, R. L. 1994. Perpetuation Theory and the Long-Term Effects of School Desegregation. *Review of Educational Research*, 64(4).

Whitehurt, G. J. 2017. New evidence on school choice and racially segregated schools.

Zhang, H. 2008. White Flight in the Context of Education: Evidence from South Carolina. *Journal of Geography*, 107(6): 236–245.