

Optimization Meets Participation: Iterative Zone Generation for School Assignment

Mobin Yahyazadeh Jeloudar, Katherine Mentzer, Irene Lo

Department of Management Science and Engineering, Stanford University,
yahyazad@stanford.edu, kmentzer@stanford.edu, ilo@stanford.edu

Abstract: Many public school districts in US draw neighborhood school boundaries that assign students to schools, based on their location proximity. Such boundaries have been criticized for creating homogeneous communities, limiting diversity, and reinforcing socioeconomic disparities. Research shows that the socioeconomic status of a neighborhood is a significant predictor of neighborhood's academic achievement, and since school districts typically align with residential neighborhoods, the traditional attendance areas can create unequal access to educational opportunities.

In this paper, we create new larger 'zones', allowing students to apply to any school within their zone. First, larger zones will provide families with more options to find a school that best suits their child's needs and interests. Second, by creating new zone boundaries that consider socioeconomic diversity and prioritize equity, we can work towards ensuring that every child has access to a high-quality education, regardless of their zip code. Diverse schools not only provide educational benefits but also promote social and emotional growth, preparing students for the realities of a diverse world. Therefore, changing the zone boundaries for public school assignments is necessary to create more inclusive and equitable schools that prepare students for a diverse society. Our work informed the design and approval of a zone-based policy in 2020 for use starting the 2026-27 school year, and is being used to support policy implementation.

1. Introduction

For U.S. public schools, geographic boundaries have long played an important role in determining which students attend which schools. For many districts, the zones take the form of neighborhood boundaries used to determine each student’s neighborhood school. In a recent student assignment policy redesign, the San Francisco Unified School District passed guidelines calling for multi-school zones. Each zone contains multiple schools. Families can express preferences over schools within their home zone, and students are assigned in a centralized way to one of the schools in their zone. Multi-school zones offer a promising middle ground between neighborhood assignment and an unrestricted citywide choice process, balancing between the segregatory forces of residential segregation and choice-driven segregation. However, the problem of optimizing multi-school zones has not been well studied.

The zone design task is to design a set of geographic boundaries that induce the best zone-based student assignment policy. However, this problem is both computationally and socially difficult. Zones impact important outcomes such as school diversity directly through geographic restriction, and also indirectly through the interaction of such restrictions with families’ choice patterns. Moreover, the zone design problem is highly constrained. Imposing ideal constraints often leads to infeasibility, requiring a human-centered approach to identify relaxable constraints, despite potential disagreements among stakeholders.

We propose an iterative framework to address the computationally challenging and multi-objective zone design optimization problem. Our framework consists of two key components: 1) a Consumer Search process that uses interactive tools to engage stakeholders in identifying and selecting zoning constraints, and 2) Computational Methods to generate solutions satisfying these constraints. In the Consumer Search process, we identify constraints that match stakeholders’ preferences. We update the constraint set by modifying metrics, adding new constraints, and adjusting constraint formats and thresholds based on stakeholder input. Our primary goal is to help stakeholders and decision-makers explore the frontier of feasible school zonings and understand the trade-offs between different solutions under various constraints and thresholds. In the Computational Methods component, we employ optimization models and Markov Chain Monte Carlo (MCMC) methods to find zone boundaries that satisfy the constraints identified through Consumer Search. We solve and accelerate these models using multi-level zone design and recursive approaches. Generating each zoning solution is computationally intensive, so we must repeatedly solve the problem with different constraint sets to support the iterative Consumer Search process. Our iterative framework

alternates between the Consumer Search and Computational Methods components, enabling stakeholders to refine their preferences based on generated solutions and feasibility constraints, thereby exploring trade-offs between different zonings.

We find that the multilevel optimization approach consistently produces better zone-level statistics compared to other methods. To evaluate the impact of zones on student outcomes, we use a simulation engine that models student preferences and assigns students to schools under each zoning policy. Comparing the resulting assignments to the current policy, we find that larger zones generally perform better on diversity metrics, with the 6-zone and 7-zone policies being the only ones that reduce both the number of high-poverty schools and the percentage of historically underserved students in those schools. Conversely, smaller zones tend to perform better on proximity measures, with the 18-zone policy substantially increasing the share of students assigned to a school within half a mile. These methods and findings will support the San Francisco Unified School District (SFUSD) in selecting zones for its new student assignment policy, scheduled for implementation in the 2026-2027 academic year.

2. The Multi-School Zone Design Problem

Problem Statement. In 2018, the San Francisco Unified School District (SFUSD) initiated an exploration to redesign their elementary student assignment policy. The district aimed to create a system balancing neighborhood assignment and citywide choice, offering students multiple school options within a limited area. The goals of the redesign were to develop a simple, predictable assignment system that would balance school choice benefits with practical distribution constraints, while strengthening community connections to local schools, reducing cross-city travel, and creating integrated schools reflecting San Francisco's diversity.

The initial proposal included geographic zones containing multiple schools, with families having access to schools within their zone. In December 2020, the SFUSD Board of Education approved the use of a zone-based plan without specifying exact boundaries. This paper outlines our work on the zone design process and the facilitation of selecting a final zone map for implementation.

2.1. Understanding the Complexities of Multi-School Zone Design

Designing multi-school attendance zones is a complex task that involves creating geographic boundaries to optimize the assignment of students to schools. This problem presents both computational and social challenges. The zones must satisfy multiple balance constraints, such as ensuring an even

distribution of student population across zones, maintaining a similar percentage of racial minority students in each zone, and balancing the population of students with low socio-economic status between zones. These competing objectives, along with the need to consider factors like compactness and contiguity, make the zone design problem computationally difficult.

While the problem of designing boundaries for neighborhood schools has been well studied (1, 2, 3), multi-school zone design introduces an additional complexity: the need to consider families' school preferences within each zone. In the neighborhood school setting, students are simply assigned to their local school, so the neighborhood composition directly translates to the school composition. In contrast, when designing multi-school zones, we must account for the fact that families have a choice among schools within their zone. This means that the ultimate goal is to impact the characteristics of students after their assignment to schools, not just the zone-level statistics. Incorporating this choice process into the zone design problem is crucial, as it influences school-level outcomes such as diversity. Abeledo and Rothblum (4) demonstrate how to implement the Deferred Acceptance algorithm with an integer program given student preferences. However, combining the zone design and choice estimation tasks into a single optimization problem leads to an extremely large optimization problem with the number of variables on the order of $O(n^4)$, making it intractable even with computational tricks and approximations (see Appendix 7.2).

In addition to the computational challenges of building zones, the multi-school zone design problem also includes a zone policy selection problem similar to problems encountered in the human-computer interaction (HCI) space. The zone optimization problem is highly constrained, and implementing all of the district's ideal constraints (i.e., exact racial balance among zones) yields no feasible zoning solutions. In the absence of computation tools, stakeholders are often unable to precisely quantify their desired constraints (e.g. how tight the balance constraints need to be), as such decisions involve implicit value trade-offs with other constraints specific to their school district (e.g. the geography of segregation, history of educational inequality). Moreover, stakeholders may disagree on which constraints should be relaxed to obtain feasible solutions, given the inherent trade-offs involved. As a result, we need to understand stakeholders' opinions, while identifying the combination of constraints that can be relaxed to yield feasible zone solutions.

2.2. Framework

To address both human-computer interaction and computational challenges, we propose an iterative framework consisting of two key components. The first component is a *Consumer Search*

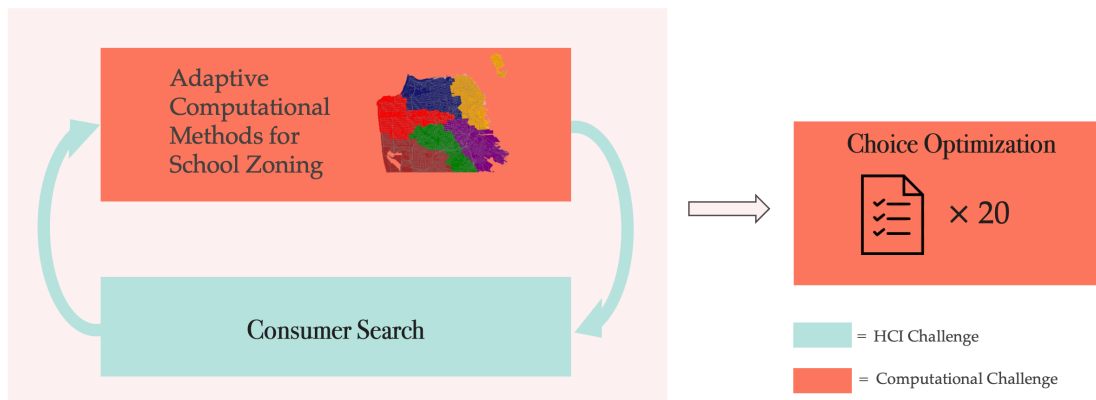


Figure 1: Iterative framework for multi-school zone design. Blue step: consumer search process to identify zoning constraints based on stakeholder preferences. Red steps: computational methods to generate zones satisfying constraints while optimizing zone shapes and balance.

process to engage stakeholders in identifying and selecting zoning constraints. The second component, *Computational Methods for School Zoning*, generates solutions that satisfy these constraints while optimizing zone shapes and balance.

Figure 1 depicts the framework, with computational challenges in red and human-computer interaction components in blue. Stakeholders refine their preferences through an iterative process of identifying constraints and generating corresponding zoning solutions. Section 4 describes an interactive tool we built using large language models to help combine these steps. Section 3 provides a comprehensive description of the various computational approaches employed to create school zones that meet the identified constraints while optimizing zone shapes and satisfying additional balance requirements when feasible.

Since the full problem of capturing the impact of choices while designing zones cannot be solved directly, we initially focus on generating an ensemble of zones and address choice in a later step. This ensemble approach narrows the problem space, allowing for subsequent optimization over a more manageable set of solutions. To address the shortcoming of our initial model, we evaluate these zones using a choice-based assignment process. We employ a simulation engine from companion work, which incorporates a multinomial logit choice model, district priorities, and the Gale-Shapley deferred acceptance algorithm (5) with diversity reserves. The resulting assignments are then used to evaluate how well the zone-based policy meets the needs of the district.

3. Computational Methods for School Zoning

This section focuses on computational approaches for generating school zones based on the constraints identified through the Consumer Search process. We adapt techniques from school neigh-

neighborhood boundary design and political redistricting to tackle the problem of multi-school zoning, and apply our techniques in a case study from San Francisco. Our goal is to generate an ensemble of solutions that represent different design decisions or value judgements, aiming to find zones that align with those values while achieving compact shapes and satisfying additional balance constraints when possible. The constraints used in this section were identified via a Consumer Search process in close collaboration with staff at SFUSD. This helped identify the most restrictive and balanced set of constraints that result in feasible solutions.

The school zoning problem shares similarities with other geographic boundary design problems, such as political redistricting. Political redistricting aims to draw compact and representative boundaries delineating political districts. Residents of each district elect a single representative, often for the U.S. House of Representatives or state legislatures. In this context, the goal is typically to find a ‘fair’ districting, where fairness is defined by the number of elected officials of a particular group in a given map relative to the average number of representatives in all other maps.

Optimization-based methods, commonly used in both school neighborhood design and political redistricting, employ integer programming or constraint programming to assign geographic units to zones based on specific objectives (6, 7, 8, 9). While these methods align well with the multi-school boundary design problem due to their ability to incorporate hard constraints, the computational complexity of solving large-scale integer programs necessitates reducing the problem space appropriately for the given context.

In political redistricting, ensemble-based methods like Metropolis Markov Chain Monte Carlo (MCMC) are used to generate large, diverse, and representative samples of possible redistricting plans (10, 11). These methods model redistricting as a graph partition problem and provide a way to transition between neighboring partitions. The representative sampling achieved by these methods is crucial for assessing the fairness of a given map relative to the ensemble. However, the multi-school zone design problem differs from political redistricting as it prioritizes finding a set of well-performing maps rather than ensuring representativeness in the generated ensemble.

We adapt both optimization-based methods and MCMC methods to provide novel computational approaches to the school redistricting problem. Both methods construct zones by combining smaller geographic units, such as census blocks, census block groups, or attendance areas¹, depending on

¹The City of San Francisco is divided into approximately 6,000 census blocks by the US Census Bureau. These blocks are further aggregated into 579 census block groups, each containing 8 to 12 blocks. SFUSD also partitions the city into 59 attendance areas, which represent school neighborhoods under the current assignment policy. The boundaries between census block groups and attendance areas do not always align. Census blocks are the smallest geographical unit that we use to solve the zoning problem.

the desired level of granularity. In this section, we present a comprehensive description of our novel computational approaches. We first discuss the key differences between school zoning and political redistricting in 3.1. We then present optimization-based approaches for school zoning, including integer programming and constraint programming techniques in 3.2. We then examine ensemble-based MCMC methods adapted from political redistricting in 3.4. In appendix 7.7, we analyze the problem of jointly optimizing resource allocation and designing zones in a single optimization step, using the example of assigning special education students to their programs.

3.1. School Zoning vs. Political Redistricting: Computational Challenges

While school zoning and political redistricting share some similarities, there are several key differences that make school zoning a more challenging problem. These distinctions necessitate the development of tailored approaches, as existing methods from political redistricting literature cannot be directly applied.

Limited Feasible Solutions. School zoning involves a more extensive set of constraints compared to political redistricting, such as balancing student populations, racial diversity, and socio-economic status across zones. The increased number and complexity of these constraints significantly reduce the number of feasible solutions that satisfy all the desired criteria. Consequently, local search-based algorithms struggle to effectively explore the solution space, frequently encountering infeasible options.

Skewness of Population Distribution In school zoning, there are multiple metrics that need to be balanced across zones, including the shortage of seats in a zone compared to the number of students in a zone, and the number of schools in a zone. The distribution of school seats across geographical units is highly skewed, with each school acting as a geographical unit with a very high concentration of seats compared to the surrounding units. Schools create a significant spike in seat capacity that must be accounted for when balancing enrollment across zones. Additionally, the distribution of the number of schools across blocks is also highly skewed, with most blocks having zero schools and a small fraction of blocks (e.g., 0.7%) containing one school.

In contrast to school zoning, the population of voters in political redistricting tends to be more evenly distributed across the geographic units under consideration. Although there may be some variation in population density, extreme concentration of voters in specific areas is less common. This relatively uniform distribution facilitates the balancing of populations across districts by algorithms that use local search-based methods (e.g., swapping blocks on boundaries between zones) to

make incremental changes to district boundaries. However, in school redistricting, the movement of a geographical unit containing a school from one zone to another can lead to drastic changes in seat counts, resulting in significant imbalances in seat shortages or surpluses and the number of schools across zones. These imbalances are challenging to resolve through incremental adjustments, often causing local search-based algorithms to become stuck in suboptimal solutions.

3.2. Optimization-Based Methods

Optimization-based methods are commonly used in both school neighborhood design and political redistricting (3, 8, 6, 7, 9, 12, 13, 14, 15, 16). These methods typically employ integer programming to assign geographic units to zones according to specific objectives. In the case of the San Francisco Unified School District, directly solving the optimization model for the entire school district is computationally challenging. To address this, we present two optimization-based approaches adapted from existing techniques, aiming to reduce the problem size and computational burden. The first approach is a multi-level method that utilizes solutions from less granular levels (e.g., attendance area) to initialize a partial solution at more granular levels (e.g., block). The second approach employs recursive methods to handle the challenge of dividing the city into a large number of zones.

Multilevel algorithms are a class of graph-based algorithms that create multiple levels of graphs by an iterative coarsening procedure (where adjacent nodes are merged to create a smaller graph) and then mapping a district solution from the reduced graph instance back to the the original graph (17, 18, 19, 20). These algorithms are considered the most effective for graph partitioning problems in applications such as VLSI design and image processing (21). The multilevel paradigm has also been adopted for other combinatorial optimization problems, such as the traveling salesman problem (22) and graph-drawing (23). We adapt the multi-level approach to school zoning and demonstrate that solving a less granular optimization problem can provide valuable insights and serve as a starting point for generating more refined solutions.

Recursive algorithms have been employed to increase the speed of optimization in political redistricting (24, 8, 25, 26). Gurnee and Shmoys (8) propose a two-stage optimization approach, similar to our method, using a randomized divide-and-conquer column generation heuristic to optimize fairness. They first generate one trillion distinct district plans by exploiting the compositional structure of graph partitioning problems and then optimize over this set. Other recursive approaches include the shortest split-line algorithm (25) and the diminishing halves algorithm (26). We introduce an overlapping recursive approach to school zoning and show that it provides a viable

approach, particularly in settings where desired zones are of similar size with geographic units used in the multi-level approach.

Constraint programming offers an alternative to integer programming, leveraging both arithmetical and logical algebra to handle nonlinear constraints. Gillani et al. (3) apply constraint programming to redraw attendance boundaries around each elementary school, aiming to reduce school segregation, using constraint programming. We find that a constraint programming approach can provide a 3x speed-up compared to the direct mixed-integer programming approach (section 3.3).

3.2.1. Mixed Integer Programming Formulation We formulate the school zoning problem as an integer program, drawing on established approaches in the literature (13, 3). Our formulation adapts these existing models to the specific context of school zoning. Our main adaptation is incorporating pre-defined zone centroids to help enforce contiguity, as well as an alternative method that does not rely on predefined centroids. Let U denote the set of predefined geographic units and Z represent the set of zones, each associated with a centroid unit $z \in U$. For each unit $u \in U$ and zone $z \in Z$, let $x_{u,z} \in 0, 1$ be a binary decision variable indicating whether unit u is assigned to zone z . The integer program aims to allocate geographic units to zones while adhering to the following objectives and constraints:

Feasibility The following constraint ensures that every unit belongs to exactly one zone:

$$\sum_{z \in Z} x_{u,z} = 1 \quad \forall u \in U \quad (1)$$

Contiguity. Contiguity ensures that every part of a school zone is reachable from any other part without leaving the zone, which is important for school commutes and interpretability. However, incorporating contiguity constraints into an optimization approach is challenging. We adapt a contiguity constraint formulation from (13), which requires pre-defined zone centroids as starting points for building each zone. Let $d_{u,z}$ be the distance between unit u and the centroid of zone z .² The contiguity constraint is given by:

$$x_{u,z} \leq \sum_{v: v \in N(u), d_{v,z} \leq d_{u,z}} x_{v,z} \quad \forall u \in U, z \in Z \quad (2)$$

Equation (2) ensures that unit u is assigned to zone z only if at least one of its neighbors that is closer to the zone centroid is also assigned to the same zone. This constraint guarantees a “path” of closer neighboring areas connecting each unit to the zone centroid.

²Distance is calculated using the geographical center of each unit, based on the 2020 census shapefile.

While constraint (2) is sufficient for contiguity, it is not necessary, and may exclude some valid contiguous zone maps (Appendix 7.4). Alternative approaches to enforce contiguity include using different distance measures, such as graph distance³, solving without contiguity constraints (27) then using a heuristic to modify the discontinuous solution (28), employing complex constraints on the units' dual graph (29)⁴, or using cutting plane methods (30).

In our problem, a centroid is a pre-defined geographical unit within each zone that serves as a reference point for ensuring contiguity. Unlike attendance area redistricting (3, 2) where each zone contains one school as a natural centroid, our zones consist of multiple schools, complicating centroid selection. We address this by considering high-popularity schools dispersed throughout the city as centroids and building zones around them. In collaboration with SFUSD staff, we generated 90 different sets of centroids to explore a wide range of solutions. While an alternative approach could treat centroids as variables in the optimization model, this substantially increases computational complexity, making it intractable to solve in reasonable time.⁵

Compactness. A common preference shared by SFUSD officials and families in school redistricting is for *compact* zones. Compactness requires zones to have a reasonable shape, avoiding unnecessary elongation or twisting. Compact zones provide a safer environment for children walking or biking to school, improve traffic flow, reduce congestion, and foster a sense of community. Moreover, contiguous and compact zone maps are simpler and easier for students and families to understand.

There are several ways to quantify district compactness, such as summing district perimeters, comparing a district's area and perimeter, summing the distance from each geographic unit to the center of its district, or counting the number of cut edges in the graph representation (i.e., edges

³Our results are shown using Euclidean distance. We also implemented contiguity with driving distance and graph distance as well, but did not pursue this approach as it significantly restricted the number of contiguous zones compared to Euclidean distance.

⁴King et al (29) enforce the following constrains on the dual graph of the set of units, for all centroids z and sets $S \subseteq (U - N(z) - z)$.

$$\sum_{u \in \bigcup_{S \in S} N(S) - S} x_{u,z} - \sum_{v \in S} x_{v,z} \geq 1 - |S|$$

This approach requires a number of inequalities that grows exponentially with the number of units, making it intractable for large problems. To overcome this, they propose running a breadth-first search on the "Geo-Graph" which they show to be more efficient than some alternatives.

⁵Treating centroids as variables would eliminate the need for predefined centroids and provide more flexibility in zone formation. However, it significantly increases the number of variables, making the problem computationally intractable using standard solvers like Gurobi.

whose endpoints are in different districts) (15, 31, 32, 10, 33).⁶ In the student assignment setting, we encourage compactness by minimizing the number of cut edges.

$$Obj_{IP} = \min \sum_{u \in U, v \in N(u)} b_{u,v} \quad (3)$$

$$b_{u,v} \geq |x_{u,z} - x_{v,z}| \quad \forall u \in U, v \in N(u), z \in Z \quad (4)$$

The objective function in (3) minimizes the total number of cuts between neighboring areas that are assigned to different zones.⁷ For each unit $u \in U$, the set $N(u)$ consists of its neighboring units.⁸ The decision variable $b_{u,v} \in 0, 1$, referred to as the boundary cost, indicates whether units u and v are assigned to different zones. Constraint (4) ensures that $b_{u,v}$ takes a value of 1 if and only if the adjacent units u and v are assigned to different zones.

Balance Constraints. Another set of constraints common in school redistricting are balance constraints within and between zones. School zones should be of a similar size, and provide sufficient seats for students in the zone. In addition, to guide toward goals of diversity and equity, it is often desirable to require some level of demographic similarity between zones.⁹ We capture these balance constraints as follows.

$$1 \geq \left| \sum_{u \in U} \frac{Sch_u}{|Z|} - \sum_{u \in U} Sch_u \cdot x_{u,z} \right| \quad \forall z \in Z \quad (5)$$

$$\alpha_{shortage} \sum_{u \in U} P_u \cdot x_{u,z} \geq \left| \sum_{u \in U} (P_u - Seat_u) \cdot x_{u,z} \right| \quad \forall z \in Z \quad (6)$$

$$\sum_{u \in U} f_u \cdot x_{u,z} \geq \left(\frac{F}{N} - \alpha_{SES} \right) \cdot \sum_{u \in U} P_u \cdot x_{u,z} \quad \forall z \in Z \quad (7)$$

⁶ Duchin and Tenner in (32) do a survey on several contour based scores including Polsby-Popper, Schwartzberg, Reock, Population Polygon and Minimum Convex Polygon score. and provide important short-comings of each scores including Coastline effects, Resolution instability, Coordinate dependence and Empty space effects.

⁷ We initially considered incorporating various balance metrics into the objective function alongside compactness. However, this led to stakeholder's confusion about the relative importance of each metric and the reasoning behind the suggested solutions. To improve transparency and control over trade-offs, we decided to focus solely on compactness in the objective function and manage other balance metrics through constraint levers. After experimenting with different compactness measures, including combinations of cut edges and squared distances from the centroids, we found that using cut edges alone yielded the best results for our use cases, promoting compact and contiguous zones while reducing computational complexity.

⁸ $N(u)$ is the set of units that share a border with unit u according to the 2020 census map shapefile.

⁹ In prior work on redistricting for attendance area / neighborhood schools, the composition of the zone is the same as the composition of the schools, and thus is a direct measure of diversity. In SFUSD, while choice allows post-assignment measures of diversity and equity to differ from zone-level ones, district stakeholders still believed it was important to aim for some level of demographic similarity between zones (8, 34, 13, 2).

$$\sum_{u \in U} R_u^k \cdot x_{u,z} \geq \left(\frac{R^k}{N} - \alpha_k \right) \cdot \sum_{u \in U} P_u \cdot x_{u,z} \quad \forall z \quad (8)$$

$$\begin{aligned} & \in Z, k \\ & \in K \end{aligned}$$

Zone size. For each $u \in U$, let Sch_u be the number of schools in unit u . Constraint (5) ensures all zones have the same number of schools (up to 1 school).¹⁰

Seat Shortage. For each $u \in U$, let P_u and Seat_u be the population of students and number of seats in unit u respectively. Constraint (6) captures the difference between the number of students and number of seats in each zone, and ensures that no zone has a seat shortage that is α_{shortage} more than its total student population. In SFUSD, we set α_{shortage} a varying number between 15% and 25%.¹¹

Demographic Parity. For each $u \in U$, let f_u be the total socio-economic score of students in unit u , and let $F := \sum_{u \in U} f_u$ be the total score over all units.¹² Similarly, for a set of races K ¹³, and for each unit u and race k , let R_u^k be the number of students of race k in unit u , and let $R^k := \sum_{u \in U} R_u^k$ be the total number of students of race k . Let $N := \sum_{u \in U} P_u$ be the total number of students in all units. Constraints (7) and (8) ensure that the average socioeconomic score of students in each zone is within α_{SES} of the average socioeconomic score of students in the whole district, and the proportion of students of a given race k in each zone is within α_k of the proportion in the district. In the SFUSD setting, we set α_{SES} and α_k to be a value between 12% and 15%, depending on the size of the zones being generated.¹⁴

¹⁰ Unlike in political redistricting, where the principle of ‘one person one vote’ is constitutionally mandated, the population count does not need to be exactly balanced between school zones. In SFUSD, schools have programs of vastly different sizes (e.g. 22 students vs. 100+ students), and district decision-makers deemed that it was more important to balance the number of options within zones rather than the number of students within zones.

¹¹ Balancing school seats and student population perfectly is challenging due to year-to-year variations in student population across zones. Moreover, about 20% of students who apply to SFUSD schools eventually do not enroll (opting for other schooling options), which helps alleviate the shortage. We vary the allowed shortage based on zone size: for larger zones with 10 schools each, we set $\alpha_{\text{shortage}} = 15\%$, while for smaller zones with 3 schools each, we use $\alpha_{\text{shortage}} = 25\%$, as it is easier to satisfy balance constraints more tightly with larger zones.

¹² In line with the literature (35, 36, 37), we measure socio-economic status using eligibility for Free or Reduced Priced Lunch (FRL), which is commonly used as a proxy for socioeconomic need in U.S. public schools.

¹³ In the SFUSD setting, we let $K := \{\text{White, Asian, Latinx}\}$, as these are the largest ethnicity groups in San Francisco. Other racial groups, such as African American students, comprise only about 5% of the total student population, making it challenging to balance them across zones.

¹⁴ For larger zones, we use stricter balance constraints (i.e., lower values of α_{SES} and α_k) as it is easier to achieve socio-economic and racial balance when the zones are more expansive. As we generate smaller zones, we gradually relax these constraints (i.e., allow higher values of α_{SES} and α_k) to account for the increased difficulty in maintaining perfect balance with more compact zones.

3.2.2. Multilevel Zone Design via Progressive Refinement Our goal is to solve the MIP formulation of the problem using census blocks as geographic units. This results in 23,847 variables and 116,645 constraints, which is computationally intractable. To address this computational challenge, we employ a multilevel optimization procedure (Figure 2) that progressively refines the solution from less to more granular geographic units, adapting the optimization parameters to the SFUSD context at each level, as detailed below.

Algorithm 1 Multilevel Zone Design

Initialization Step: Solve the mixed integer programming optimization problem in section 3.2 using attendance areas as geographic units. (Figure 2a)

Iterative Optimization: Iteratively use the solution from the previous step, at a lower granularity level, to find a solution at a higher granularity level.

- **Mapping and Trimming Step:** Initialize a partial solution using the next level of geographic granularity.¹⁵ Map each unit to the zone of its corresponding lower-granularity unit from the previous step. Trim the partial solution by discarding the zone assignment for any units with only a few adjacent units from the same zone (Figure 2b and 2d). Tighten the balance constraints by about 8% compared to the previous step to achieve better overall balance in the final zoning solution.¹⁶
- **MIP With Initialized Variables:** Initialize the MIP variables based on the assignments from the previous step (at either the block group or block level) and solve for the remaining variables, ensuring the constraints from subsection 3.2 are satisfied (Figure 2c and 2e).

Local Search Step: Improve the solution obtained at the block level using a local search algorithm that explores the solution space by randomly selecting two adjacent zones, merging them, and attempting to cut them into two zones again (Figure 2d). Details presented in Appendix 7.5.

¹⁵ If the previous step solved the problem using attendance areas, initialize a partial solution using block groups. If the previous step used block groups, initialize a partial solution using census blocks.

¹⁶ For instance, if the goal is to find school zones with a maximum SES deviation of 20% at the attendance area level, the model will aim for 18% SES deviation at the block group level and 16.5% SES deviation at the block level. This approach allows for satisfying more restrictive constraints and achieving better overall balance in the final zoning solution by leveraging the ability to find more balanced solutions when working with smaller units.

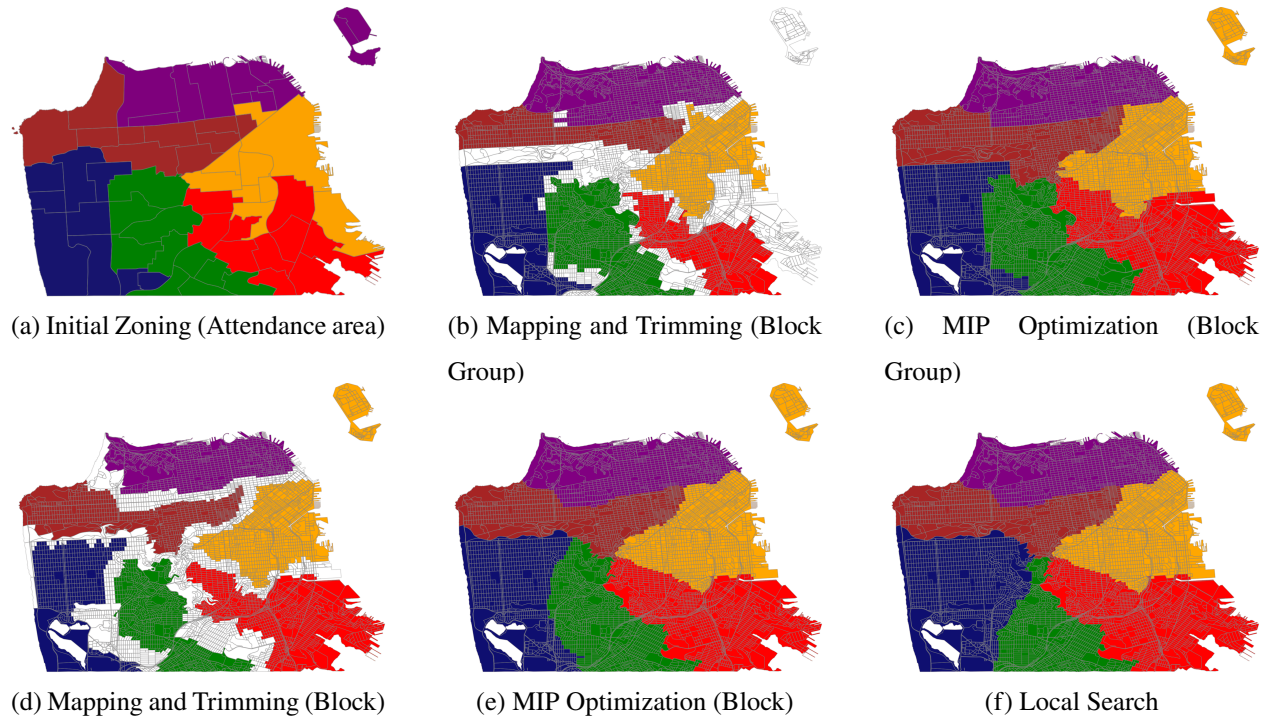


Figure 2: The 6 steps of the multilevel zone design.

The trimming step promotes zone compactness by removing isolated units, enabling the optimization model to find improved solutions at the more granular level. While units near the zone centroid remain in the same zone, boundaries can vary considerably for units further away. The smaller geographical units in the later stages provide increased flexibility, allowing the model to create more balanced zones, which is why the balance constraints can be tightened at each step of the multilevel procedure.

3.2.3. Overlapping Recursive Zoning To enhance computational efficiency at the census block level, we also introduce an overlapping recursive algorithm that sequentially solves subproblems by dividing the district into overlapping subsections.

The integer programming model for Overlapping Recursive Zoning is similar to the previously described model [3.2.1](#), with most constraints imposed in the same way. The objective is to minimize the boundary cost, and we enforce shortage/overage constraints to ensure that each zone does not exceed the shortage or overage thresholds. The primary difference in the recursive approach is how it handles feasibility constraints. Specifically, instead of recursively partitioning the district, it divides

the district into overlapping subsections, and enforces extended feasibility constraints on the regions of overlap. [17](#).

Our algorithm iteratively optimizes zones around specific centroid sets. Let S_{opt} represent the centroids targeted for zone optimization in a given subproblem (e.g., 7 schools in western San Francisco, Figure [3a](#)). Additionally, let $S_{extended}$ denote an expanded set of centroids that encompasses S_{opt} , providing a buffer for the recursive optimization algorithm (e.g., 14 schools in western San Francisco, Figure [3d](#)).

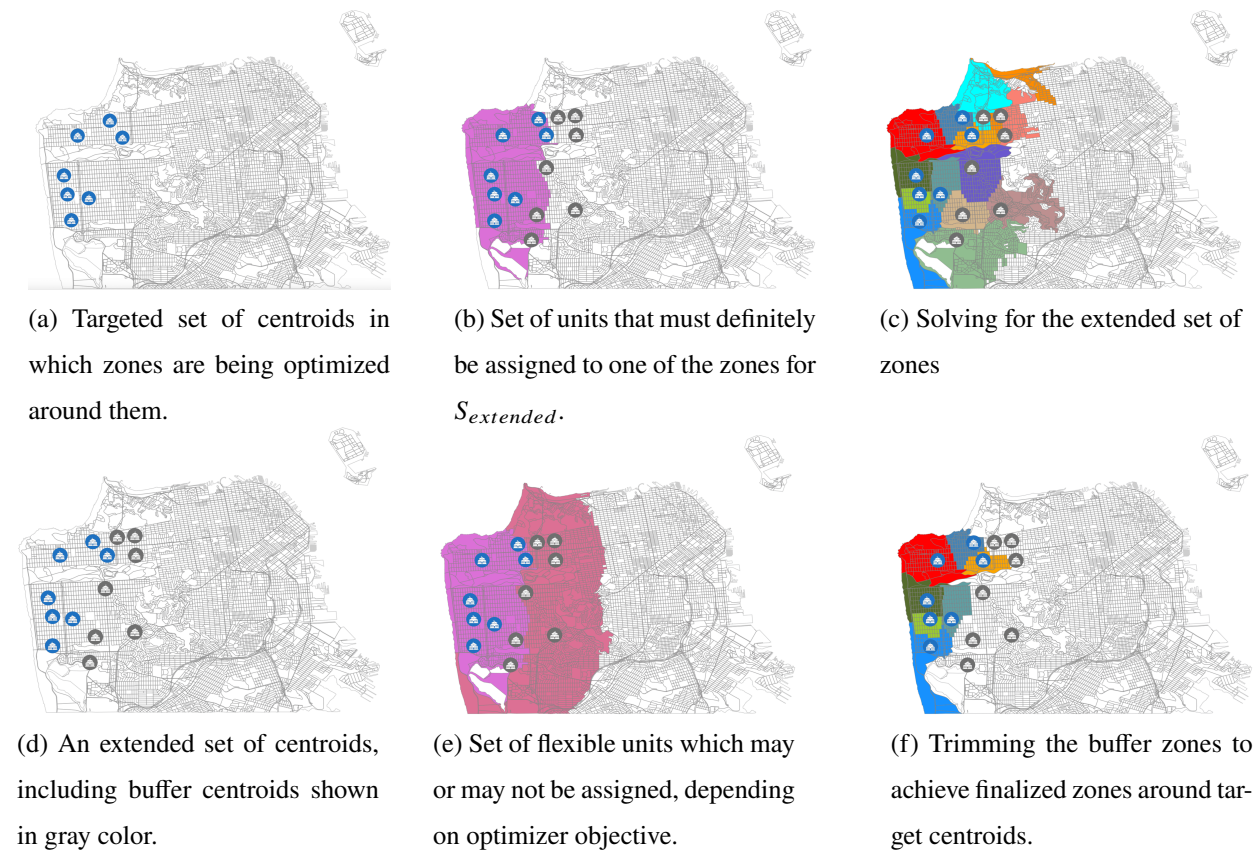


Figure 3

Extended Feasibility Let U_{must} be a small set of units within a certain distance (e.g., 3 miles) from centroids in S_{opt} that must be assigned to zones with centroids in $S_{extended}$ (Equation [9](#), Figure [3b](#)). Let U_{flex} be a set of flexible units that may or may not be assigned to zones with centroids in $S_{extended}$ (Equation [10](#), Figure [3e](#)). Both U_{must} and U_{flex} are determined based on the targeted set of centroids. The remaining units should not be assigned to any centroids in $S_{extended}$.

¹⁷The feasibility constraint determines which units should be assigned to the subproblem zones being optimized. Dividing the district into two sections (e.g., *West* and *East*) and enforcing a feasibility constraint $\sum_z x_{u,z} = 1$ for all $u \in West$ would result in the same issues as the Divide and Conquer algorithm [7.6](#)

$$\sum_z x_{u,z} = 1 \quad \forall u \in U_{must} \quad (9)$$

$$\sum_z x_{u,z} \leq 1 \quad \forall u \in U_{flex} \quad (10)$$

$$\sum_z x_{u,z} = 0 \quad \forall u \in U \setminus U_{must} \setminus U_{flex} \quad (11)$$

Equation 10 allows the optimizer to decide if assigning U_{flex} units to $S_{extended}$ centroids improves the objective function while satisfying all constraints. Equation 11 ensures that units more than a specific distance (e.g., 5 miles) from all $S_{extended}$ centroids are not assigned to any $S_{extended}$ school.

Algorithm 2 Overlapping Recursive

Iteratively solve subproblems:

- Select $S_{extended}$ and S_{opt}
 - Select U_{must} and U_{flex}
 - Solve the optimization for the given set of variables and their assignment to centroids in $S_{extended}$ (Figure 3c)
 - Fix and finalize the zones around centroids in S_{opt} (Figure 3f).
 - Proceeds to solve the next subproblem.¹⁸
-

The overlapping recursive algorithm offers greater flexibility in defining feasibility constraints compared to the Divide and Conquer approach. It allows the optimizer to consider a larger set of units that could be assigned to the targeted zones. Moreover, the overlapping nature of subproblems help maintain compactness, as the extended school sets ($S_{extended}$) provide a buffer, enabling smooth transitions between adjacent subproblems.

By carefully selecting the target and extended school sets and determining an appropriate order for solving the subproblems, we successfully constructed the solution presented in the results section (Figure 11). This process relies on a combination of domain knowledge and experimentation to balance computational efficiency and solution quality. However, determining the appropriate size and composition of the sets $S_{extended}$, S_{opt} , U_{must} , and U_{flex} can be challenging. If the subproblems are not carefully designed to minimize overlap and redundancy, this method may result in longer computation times compared to Divide and Conquer approaches.

¹⁸For instance, after solving for the zones in the western part of San Francisco, the algorithm might continue by finding a zoning solution for the southwestern part of the district while keeping the previously optimized zones fixed.

3.3. Constraint Programming

Constraint Programming (CP) is a powerful paradigm for solving combinatorial problems, using both arithmetical and logical algebra to handle highly disjunctive (non-linear) constraints. CP is applied to a wide range of domains, including scheduling, planning, vehicle routing (38, 39, 40, 3, 41, 42, 43). In our project, we implemented the zone optimization model (Section 3.2.1) using CP, which outperformed the Integer Programming IP implementation by approximately 3 times for smaller subproblems. However, despite this speedup, CP still could not directly solve the zoning problem due to the large scale. To achieve solutions at higher granularity, we relied on speed-up techniques such as the multilevel approach (Section 3.2.2) and the recursive approach (Section 3.2.3).

3.4. MCMC Approach

The school zoning problem shares similarities with political redistricting, as both involve partitioning a geographic area into regions while balancing various criteria and constraints. MCMC methods have been effective in addressing the computational challenges in political redistricting via an ensemble-based approach. In this section, we apply this approach to generate an ensemble of school zones, adapting MCMC methods to the school zoning context.

MCMC methods have proven to be powerful tools for sampling from the vast space¹⁹ of valid redistricting plans and detecting gerrymandering (44, 45, 46, 47, 48, 49). They are particularly useful for generating large, diverse ensembles of plans that serve as baselines for comparison (10, 50, 51). By contextualizing a proposed plan within a set of alternative valid plans, MCMC can reveal whether it is an outlier in terms of partisan bias or other metrics. This ensemble-based approach has been successfully employed in legal challenges to partisan gerrymanders at both state and federal levels (10, 50, 10) and used as evidence in several recent court cases (52, 53, 54, 55).

However, the goal in school zoning differs from political redistricting. Given a set of constraints identified through constraint discovery, the aim is to find compact solutions satisfying all constraints in the set. In addition, in school zoning it is desirable to generate an ensemble of well-performing zoning plans to illustrate the trade-offs between constraints. To adapt MCMC methods to the school zoning context, we first generate an ensemble of candidate plans using adaptations of the Recombination (ReCom) and Relaxed ReCom algorithms introduced by (10) and (56). We weight

¹⁹ Even for modestly sized problems, the number of valid districting plans satisfying typical legal criteria can exceed 10^{100} (10)

the Relaxed ReCom algorithms to encourage exploration of well-performing plans. We then filter the ensemble to extract plans that satisfy the specific school zoning constraints or achieve the best balance across all criteria.

Sampling Methods for Redistricting (Flip vs. ReCom) A crucial aspect of applying MCMC methods is the choice of the target distribution from which to sample. The two primary families of random walks used for sampling in the redistricting literature are *Flip* walks (57) and Recombination (*ReCom*) walks (10).

Flip Walks operates by randomly selecting a boundary node at each step and reassigning it to one of the adjacent zones, while preserving the contiguity of the zones. This naive approach samples from the uniform distribution over all connected partitions, satisfying the population constraints. Although simple to implement and computationally fast, Flip walks suffer from slow mixing times and a tendency to generate highly non-compact zones (10).

ReCom Walks was introduced by (10) to address the limitations of Flip walks, by making more substantial changes to the zoning plan at each step. ReCom assigns weights to each zoning plan proportional to the product of the number of spanning trees within each of its zones. The spanning tree distribution has the desirable property of concentrating probability mass on more compact zones, as compact zones tend to have a larger number of spanning trees compared to irregularly-shaped zones (58, 59). Recent studies have successfully applied ReCom to redistricting problems, demonstrating its effectiveness in efficiently exploring the space of valid plans and redrawing electoral district boundaries in various states (33, 60).

Temperature-based variants of Flip and ReCom approach, such as simulated annealing²⁰ and parallel tempering, have been proposed to address the mixing time challenges encountered by these algorithms (10, 63, 47, 61, 62, 44, 48). However, as demonstrated in our results, the effectiveness of these techniques is limited when applied to the school zoning problem.

Other MCMC approaches have been proposed to address the limitations of traditional methods including Forest ReCom (64, 51), lifting (65, 66, 67) and Sequential Monte Carlo (68). The spanning tree distribution and its c-biased variants introduced by Charikar et al. (56) offer a promising starting point to design the Markov chain in school zoning. However, school zoning’s multiple balance constraints require tailored sampling approaches, which we explore next.

²⁰ Simulated annealing is a technique that has been applied to various combinatorial problems, including redistricting (47, 61, 62, 44, 48). It allows the algorithm to escape local optima by occasionally accepting worse solutions, especially in the early stages when the “temperature” is high. As the temperature decreases, the algorithm becomes more selective, favoring better solutions and converging towards the global optimum.

3.4.1. MCMC approaches for School Zoning To address the limitations of traditional MCMC approaches in school redistricting, we implement Relaxed ReCom, a variant of the ReCom algorithm (56, 10), and define a tailored sampling distribution for the school zoning context.

School zoning and political redistricting can be modeled as a graph partition problem (32, 69, 70). The problem is represented by a graph $G = (U, E)$, where U corresponds to geographical units (e.g., census blocks) and E connects adjacent units. A zoning (Z_1, \dots, Z_k) , assigns each unit $u \in U$ to one of k zones, ensuring each zone forms a connected subgraph of G . The key steps of ReCom are: (I) Randomly select two adjacent zones. (II) Form a random spanning tree of the induced subgraph of the selected zones.²¹ (III) Cut a uniformly random edge to create two new zones. (IV) If the new zones meet balance constraints, update the zoning plan. Relaxed ReCom modifies steps (III) and (IV) by relaxing balance constraints in (IV) and changing the cut edge distribution in (III) to balance constraint enforcement and exploration. While more computationally intensive per step than ReCom or Flip, Relaxed ReCom offers advantages in rapid mixing²² and distributional properties. Algorithm 3 outlines the Relaxed ReCom process.

Algorithm 3 Relaxed ReCom (56)

Input: Graph $G = (U, E)$. Zoning (Z_1, \dots, Z_k) , and the spanning tree for each zone, (T_1, \dots, T_k) .

Output: A new zoning $(Z_1^{new}, \dots, Z_k^{new})$. Spanning tree for each zone, $(T_1^{new}, \dots, T_k^{new})$.

For iteration $t = 1, 2, \dots$:

- Add a uniformly random edge e that connects two adjacent zones Z_i and Z_j .
- Let T_i and T_j be spanning trees of $G[Z_i]$ and $G[Z_j]$ respectively. $T_i \cup T_j \cup \{e\}$ forms a spanning tree for $G[Z_i \cup Z_j]$.

- **Cutting Edge:**

- For each edge f in $T_i \cup T_j \cup \{e\}$, removing f creates:

- * Two new trees T'_i and T'_j . Let Z'_i and Z'_j be the set of vertices of T'_i and T'_j .

- * Let $Z_s^f := \begin{cases} Z'_s & \text{if } s = i \text{ or } j \\ Z_s & \end{cases}$

- **Cutting Edge Selection:** Choose²³ edge f with probability proportional to a predefined distribution over the resulting zoning (Z_1^f, \dots, Z_k^f) .

²¹ Using Wilson's algorithm (71)

²² (56) demonstrates rapid mixing for grids with balanced homogeneous weights, though extending these guarantees to more general graphs remains an open question.

To implement Relaxed ReCom, we define a sampling distribution μ^b (Biased Spanning Tree Distribution) over all possible school zonings. We then employ the down-up walk method from Anari et al. (72) to sample the new cutting edge f at each step. This approach provides a mechanism for sampling edges based on any given input distribution over graph partitions.

Biased Spanning Tree Distribution: Given a graph $G = (U, E)$ and number of zones k , the school zoning spanning tree distribution is defined by,

$$\mu^b(Z_1, \dots, Z_k) = \prod_{i=1}^k T(Z_i) \cdot |Sch_i|^{w_{Sch}} \cdot |Z_i|^{w_{size}} \cdot |Shortage_i|^{w_{Shortage}} \cdot |FRL_i|^{w_{FRL}}$$

where $Z_1 \cup \dots \cup Z_k = U$ is a zoning of U , $T(Z_i)$ is the number of spanning trees in $G[Z_i]$, Sch_i is the number of schools in zone i , $|Z_i|$ is the number of vertices in zone i , $Shortage_i$ is the proportional shortage in zone i , FRL_i is the proportion of students eligible for Free or Reduced Priced Lunch in zone i , and $w_{Sch}, w_{size}, w_{Shortage}, w_{FRL} \geq 0$ are predefined fixed weights for balance in the distribution of number of schools per zone, size of each zone, shortage of different zones, and number of students eligible for free or reduced priced lunch across zones, respectively.

If we set all weights w to zero, the biased spanning tree distribution reduces to the spanning tree distribution. The term $\prod_{i=1}^k |Sch_i|$ increases as the number of schools across zones becomes more balanced, and similar behavior is observed for zone size, shortage, and FRL. Consequently, the biased distribution assigns the highest probability to zones with equal numbers of schools, size, shortage, and FRL. As w_{size} tends to infinity, the biased spanning tree distribution approximates the balanced spanning tree distribution, effectively imposing a “soft” constraint on the balance of zone sizes. Analogous arguments hold for the number of schools, shortage, and FRL in each zone.

4. Consumer Search

In this section, we present our approach to the Consumer Search component of our iterative framework. We draw on ideas from consumer search and human-computer interaction to facilitate decision-making about desired constraints and trade-offs that are achievable with zone maps. Our goal is to provide stakeholders with sufficient support to comprehend and specify zone requirements, despite it being a computationally challenging, multi-objective, and multi-stakeholder problem. Our approach to the Consumer Search problem was used iteratively with our computational methods to identify constraint sets and well-performing zone plans within those constraint sets for SFUSD.

The zone design problem is one of many applied optimization problems that face challenges in human-computer interaction (HCI). Specifically, the complexity of the problem and computational intractability make it challenging to ensure that algorithm designers accurately translate the human problem into a technical solution. Transparency in the model design process and effective communication between stakeholders and engineers are crucial for developing better solutions and preventing unintentional policy decisions (73, 74). While algorithms can improve decision-making in data-rich environments, a hybrid approach incorporating both data-driven and socially-driven factors is often more appropriate, particularly when the true problem requirements cannot be fully captured within the algorithmic formulation (75, 76).

The multi-school zone design problem involves numerous stakeholders who are heavily invested in the policy outcomes. One way to incorporate the preferences of these stakeholders is through participatory design. This well-established concept in human-computer interaction literature emphasizes end-user involvement, enabling non-expert stakeholders to provide direct input on technology design (77, 78, 79, 80, 81). Participatory methods align with democratic ideals of public policy decision-making (82, 83, 84), leading to increased open-mindedness and satisfaction (85), as well as improved outcomes in healthcare (86). In addition to participatory design, we employ Participatory Action Research (PAR), which engages stakeholders as co-inquirers in constructing research plans and interventions (87, 88, 89, 90, 91).

4.1. Consumer Search Framework.

While the framework presented in Section 2.2 suggests a single Consumer Search process, our close collaboration with the district revealed the need to engage human decision-makers before every computational step. Thus the HCI component of our framework involves three main stages: 1) identifying metrics that align with stakeholders' preferences, 2) forming a consideration set, where stakeholders sequentially screen and search for feasible constraints, and 3) selecting the set of constraints that best align with the district's objectives from the consideration set. Figure 4 provides a detailed illustration of our framework.

In the metric elicitation stage, we use methods from participatory design to engage with a variety of stakeholders to understand the metrics to use in constraints to best capture their needs and preferences. ²⁴ SFUSD staff organized several focus groups across diverse groups of stakeholders

²⁴ Our starting point was the metrics and sets of constraints used to develop the 2020 Board policy, and a set of well-performing zone maps developed to support that policy (92).

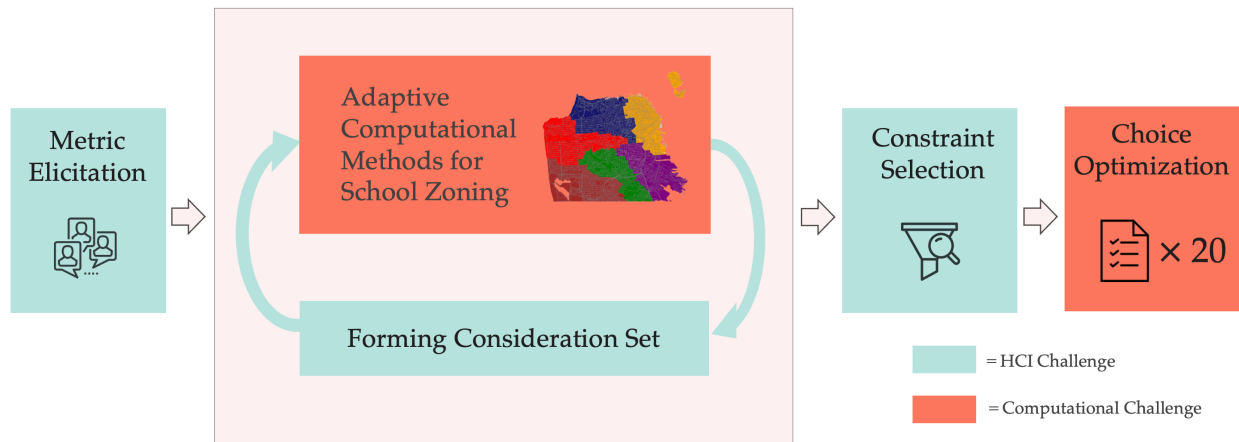


Figure 4: A detailed framework for participatory school zone design, consisting of metric elicitation, consideration set formation, and constraint selection, with stakeholder engagement throughout the process.

to gather feedback to refine our metrics, introduce new constraints, and adjust existing ones. We also developed an interactive sensemaking tool that allowed participants to interact dynamically with zone maps. This tool supports a human-centered approach to metric elicitation by providing an interactive visualization tool that makes abstract metrics concrete through maps and visuals.

In the consideration set formation stage, stakeholders sequentially submit queries to the computation tool to determine whether different sets of constraints are feasible. This allows the stakeholders to build a Pareto frontier of constraint sets that are achievable via different zone maps. This process is necessary because the search space often exceeds stakeholders' ability to effectively determine their optimal trade-off between constraint thresholds. The a priori formation of a consideration set mirrors the type of consumer behavior described in the consider-then-choose literature (93, 94, 95, 96, 97).

Once the consideration set is formed, stakeholders proceed to the selection stage, where stakeholders analyze zone maps near the Pareto frontier of constraint sets using a dashboard, and take into account additional information such as accessibility to public transportation. This tool allowed them to choose the set of constraints within the consideration set that best align with the district's objectives.

4.2. Metric Elicitation and Participatory Design (PD)

Participatory Design (PD) offers a framework for creating a hybrid space where diverse stakeholders can meaningfully contribute to the design process. We employed principles from PD to identify

the constraints and metrics crucial for zone design, bridging the gap between stakeholders' and designers' domains.

To implement this approach, SFUSD organized several focus groups, an effective method for involving a broad range of stakeholders. These groups included participants from various organizations and stakeholder segments.²⁵ The focus groups aim to gather feedback on the metrics to be balanced, introduce new constraints, and adjust the format of existing constraints.

In response to feedback from focus groups, we developed an interactive web-based tool that serves as a sensemaking object by providing a dynamic platform for engaging with the zone design problem. This Sensemaking tool addresses the limitations of static maps by allowing participants to draw their own solutions, expand the consideration set, and alleviate concerns about predetermined zone decisions. Its interactive nature facilitates the elicitation of tacit knowledge from domain experts, as they can identify and attempt to resolve non-viable aspects of the map using their contextual knowledge. The Sensemaking tool features a user-friendly interface that includes a map of San Francisco's census block groups, a live table displaying zone demographics, tools for modifying zone boundaries, and checkboxes for selecting map overlays. Users can interact with the map by selecting regions and assigning them to different zones, with the demographic table updating in real-time to provide instant feedback on the diversity and school options within each zone (Figure 5). The tool also offers optional map overlays, such as public transportation routes, school locations, community-based organizations, and district boundaries, to provide additional geographical context for informed decision-making.²⁶

The interactive zone design tool is designed for a broad user base, including focus group participants and district staff. Researchers and district staff can employ the tool in various settings, such as interviews, focus groups, and surveys, to gather valuable stakeholder input.

4.3. Forming the Consideration Set

In the school zoning problem, each constraint aims to balance certain metrics, such as the number of seats and students, across zones. The constraints allow a certain amount of deviation from the optimal balance, and the thresholds of maximum allowable deviations can be adjusted to form

²⁵ Individual focus groups were held with the SFUSD Educational Placement Center (EPC), the San Francisco Municipal Transit Agency (SFMTA), the San Francisco Department of Children, Youth, and their Families (DCYF), SFUSD's Superintendent's Executive Cabinet, parent leaders, and former SFUSD Student Assignment staff.

²⁶ The tool uses data on Kindergarten students from multiple years, aggregated at the census block group level, including student demographics, school locations, community-based organizations, local transit lines, and supervisorial districts.

Zone Design Tool

Modify the map below to design your own zones.

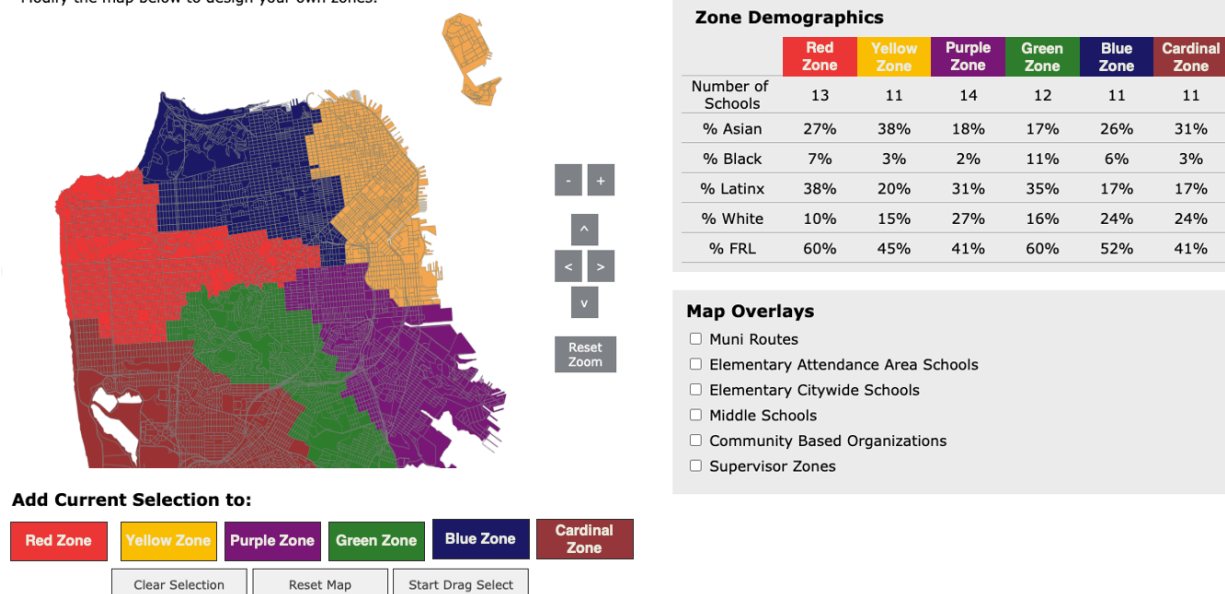


Figure 5: Interactive, web-based zone design tool's user interface.

different sets of constraints. We refer to each set of constraints and their respective thresholds (e.g., maximum allowed shortage of 20%, maximum allowed socio-economic deviation from the average of 15%) as a *position*, a term borrowed from the consumer search literature (94). Exploring various positions is crucial, as tightening some constraints may require relaxing others.

We are interested in feasible positions that are Pareto optimal, meaning no other feasible positions exist where all individual constraint thresholds are more restrictive, with at least one strictly more restrictive.²⁷ In other words, no Pareto improvement can be made without compromising the feasibility of the zoning solution.²⁸

To help stakeholders understand the range of possible zoning solutions, we consider the full range of valid options and guide them towards preferred areas of the solution space. We offer suggestions on which constraints to relax or tighten, learning the district's preferred trade-offs among design objectives. This approach parallels consumer behavior in navigating product spaces (98, 99, 100), where users continuously update their beliefs about product attributes during the search process

²⁷ Feasibility is determined by the zone generation algorithm. We use the Multilevel Optimization algorithm to generate our zones in this process. In result section 5.3, we compare the multilevel optimization with other zone generation algorithms.

²⁸ When searching for zones satisfying a set of constraints, the resulting zones often have balance deviations at the threshold values, as the optimization model's objective is set to maximize compactness. Consequently, we use the terms constraints and their corresponding zones interchangeably when discussing balance.

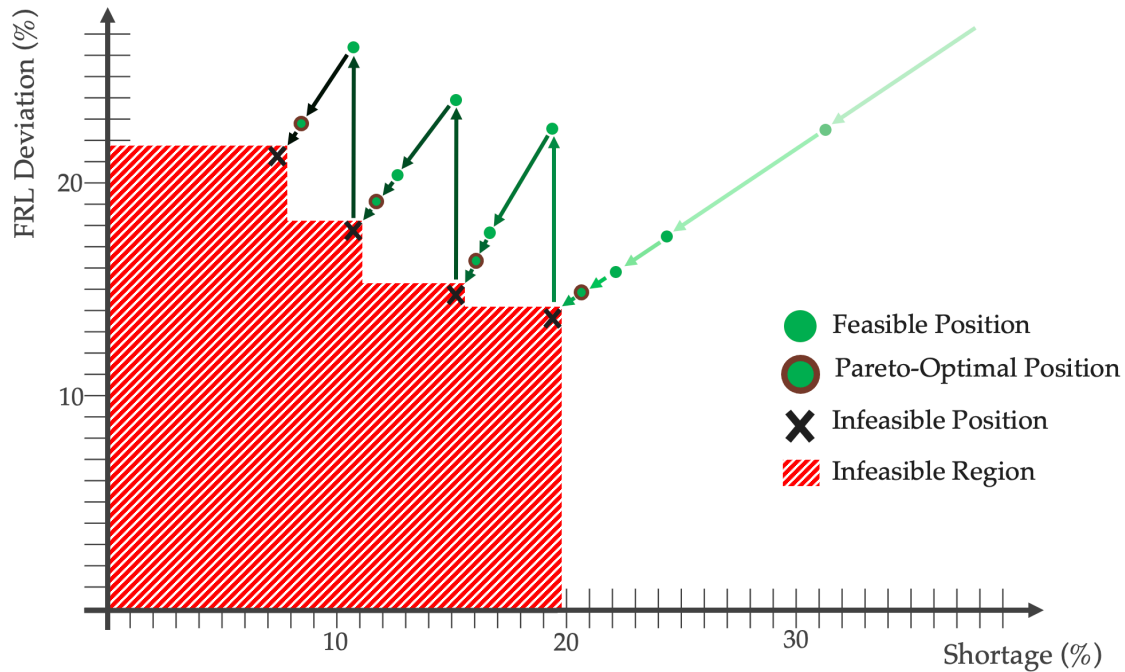


Figure 6: Stakeholders search process.

(101). We focus on determining the most lenient and most restrictive thresholds for each constraint based on stakeholder input and a multi-dimensional binary search, while still allowing for feasible zoning solutions.

Approach 1: Sequential Process Due to the large number of feasible positions, our stakeholders use a sequential screening process to build a consideration set. We begin with a set of constraints that yields a feasible solution and gradually tighten the thresholds on some constraints to achieve Pareto-improving zones until we reach an infeasible position. Upon reaching such a position, we have approximately identified a Pareto-improving zone, which is added to the stakeholder's consideration set.

Stakeholders decide whether to stop the screening and finalize their consideration set or to screen another position by querying the existence of a zoning solution for a new position. For each new query, we must relax at least one constraint to obtain a feasible solution. Through discussions with stakeholders, they decide which constraints they are willing to loosen at that position, based on the district's objectives, such as balancing socio-economic status or shortage within zones.

Figure 6 illustrates the iterative process of constraint discovery and our interactions with stakeholders. Although the figure only depicts a 2D map focusing on socio-economic balance and shortage, in practice, we consider multiple metrics, including racial distribution (e.g., the proportion of

racial minorities in each zone)²⁹ and the inclusion or exclusion of K-8 schools in the zone design. We also received feedback on new constraints, such as ensuring zones do not cross geographic barriers such as parks, mountains, and highways.

In our regular meetings with stakeholders to identify their consideration set, we begin by presenting a carefully selected set of positions that can be effectively evaluated and discussed within the available time and mental capacity of the group. At each step of modifying the constraints, we generate a range of options for potential directions to explore, with some paths leading to infeasible solutions and terminating. Figure 6 shows the trajectory of a single request. In practice, we run many requests in parallel, with each request branching into new directions at each step. We repeat this process for various zone sizes to understand the trade-offs and implications of different zone configurations.

Approach 2: Automatic Tool Using Large Language Models. To streamline the constraint exploration process and empower stakeholders to independently investigate the solution space and build their consideration set, we developed an interactive user interface (UI) tool. This web-based application allows stakeholders to directly manipulate constraint parameters or request additional constraints with the help of language models and observe the resulting zoning solutions in real-time. This significantly accelerating the iteration cycle by removing the researcher from the loop (74) and enabling stakeholders to search the feasible positions autonomously.

4.4. Constraint Selection Dashboard and Participatory Action Research (PAR)

Following the formation of the consideration set, we employed principles from Participatory Action Research (PAR) to guide stakeholder-driven constraint selection. PAR, a collaborative approach that positions stakeholders as co-researchers, was employed to collectively conduct analyses and interpret outcomes (81, 102). To support this process, we developed an interactive dashboard that enabled district staff to explore optimization-generated solutions and narrow down the consideration set from over 200 to approximately 25 zones. (Figure 16)

The dashboard consists of three selection stages, each targeting one of the district's goals: predictability, proximity, and diversity. Users navigate through tabs representing each stage. Figure 16 illustrates the diversity selection stage, where users can specify the maximum percentage of free or reduced-price lunch-eligible students within each zone and the maximum percentage of a single ethnic group relative to the entire zone population. Additional metrics capturing zone demographics are

²⁹This raises questions about which racial groups should be balanced and what level of deviation from the district average is acceptable for each group. In SFUSD, African American, Latinx, Pacific Islander students have been identified as racial minorities.

provided for context, and users can input their favorite zones and comments for later reference. The final tab offers an in-depth investigation of the user’s favorite maps. The extensive dashboard, shown in Figure 17, includes all metrics requested by stakeholders during the design process (Appendix 8.3). Users can view map visualizations and a large number of metrics describing the schools and students within each zone boundary, enabling them to make informed decisions.

5. Results

We applied our framework to generate zones and identify promising candidate maps for the San Francisco Unified School District (SFUSD). It is important to note that the selection process described is part of an ongoing zone boundary selection process and does not constitute finalized policy decisions.

5.1. Data

Each year, around 3,000-5,000 kindergarten applicants participate in the SFUSD student assignment process.³⁰ They apply to programs at the 72 elementary schools in SFUSD, including 9 K-8 programs. The K-8 programs are ‘citywide’ in the zone-based policy, meaning that all students can apply to these programs irrespective of their zone of residence.

For the zone design problem, we use the following data for kindergarten applicants from the 2014-2015 school year through the 2023-24 school year: each student’s home location (latitude and longitude), race, special education status, English language learner status, and submitted rank-order lists over SFUSD programs. Additionally, we use the 2019 free and reduced price meal eligibility rate of each census block, averaged over all grades and students within a block, as an anonymized measure of socio-economic status. We also take into account the location of schools, the capacities of each program and school-level test scores. The student and school data are then aggregated to the block, block group, and attendance area levels to characterize the geographic building blocks of zones.

Computational Resources: All runs were performed on a 2.6 GHz 6-Core Intel Core i7 with 16GB of RAM. The code was implemented in Python, with the Mixed Integer Programs written using Gurobi and the Constraint Programming approach formulated in the CP SAT-solver. Basic

³⁰ Student participation patterns have changed post-pandemic, with a drop in numbers and higher year-to-year variations compared to the pre-pandemic level of 5,000.




Figure 7: Set of Pareto Optimal positions.

model reductions were applied to reduce the number of variables and constraints when possible³¹. Optimizing directly at the block level involves 23,847 variables and 116,645 constraints.³²

5.2. Consumer Search Results

We begin with the results of the Consumer Search process. Figure 7 gives a partial view of the consideration set created by a group of stakeholders. The consideration set was obtained via the Consumer Search process, and includes a Pareto frontier of positions for zones of varying sizes (enabling 6, 10, or 18 zones in the district). Figure 7 illustrates the trade-offs between zone size, FRL balance, and shortage enabled by these positions. For example, to find a 6-zone solution with better FRL balance, higher shortage across zones is required. As expected, larger zones allow for better balance of shortage and FRL across zones. Although the figure only depicts the trade-offs between FRL balance and shortage for each zone size, in reality, there are many zone sizes, and many other metrics to balance. The finalized consideration set included 260 positions.

³¹ For example, if the distance between a geographical unit and a centroid is very large, we can be certain that the unit will never be matched to the zone corresponding to that centroid, and thus, we do not define a decision variable for that pair.

³² The computational capacity allowed for solving the direct zoning problem for approximately 5 zones and 800 units, which represented about 1/12th of the entire district.

Given the consideration set, district staff experimented with the generated solutions using the Constraint Selection dashboard described in Section 4.4. They used a scoring method³³ to choose a set of candidate zones from the consideration set. Filtering proceeded in 2 stages. First, a staff member performed a visual review of all maps, eliminating candidates that had impractical compactness failures, or zoning maps divided by highways or natural obstacles (e.g., city parks/hills). Next, maps not in the top quartile total score, or maps that were not in the top half of the diversity score, proximity score, or predictability score individually were eliminated³⁴. The results of the constraint selection process are shown in Appendix table Table 2.

5.3. Zone Generation Algorithm Comparison

After identifying a smaller set of balance constraints through the constraint selection process, we evaluate the performance of the zones generated by each of our algorithms for each of these constraints. We compare the solutions obtained when solving the optimization problem directly on attendance areas with those generated using the multilevel approach. Additionally, we compare these results to the Relaxed Recom approach.

In Figure 8, different algorithms are represented by distinct colors. The multilevel optimization consistently achieves better solutions compared to the other approaches. Figure 9 shows the zones generated by these methods, demonstrating that the multilevel approach not only results in better statistics but also provides more compact zones. Figure 9 also shows the variations of shortage and FRL metrics across zones.³⁵ The Relaxed Recom approach does not perform as well as the optimization-based methods. However, this is not entirely unexpected, given that these algorithms are not designed to find a single optimized solution.

5.3.1. Results on MCMC Approaches Figure 9c illustrates the results of applying Relaxed ReCom for 25,000 steps while attempting to balance all the relevant constraints, trying to divide the city into 6 zones. The generated solutions fail to achieve a satisfactory balance across all the desired criteria at the same time. For example, a solution that perfectly balances the number of students and population of low socio-economic status might have an uneven distribution of schools across zones.³⁶

³³ The scoring and filtering process relied on 10 working definitions: 5 diversity goals, 3 proximity goals, and 2 predictability goals.

³⁴ Details of this process are provided in Appendix 8.1.

³⁵ Our constraint ensured that the maximum shortage over all zones was within certain thresholds, but it does not provide insight into the distribution of this metric across zones.

³⁶ The zones are of drastically different sizes (e.g., the yellow and green zones have significantly different areas), and some zones, like the red zone, are overly stretched and almost comprised of two large compact areas that are far away from each other but connected

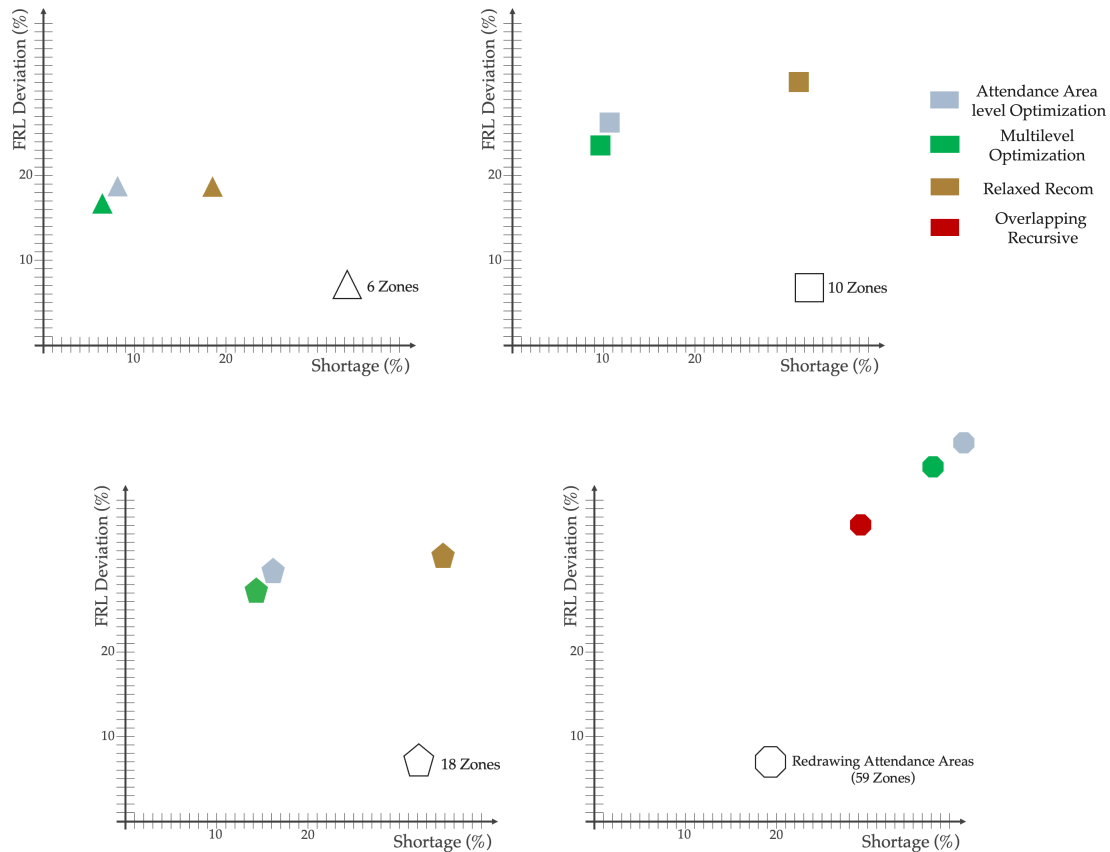


Figure 8: Comparing different zone generation methods.

Given the importance of satisfying all constraints in the school zoning context, even a failure in balancing one of them renders the entire solution inadequate for practical use. The challenges encountered by the Relaxed ReCom approach in this setting can be attributed to several factors. The large number of constraints involved in the school zoning problem creates a *high-dimensional* search space, making it harder for the algorithm to find a solution that balances among all the criteria. The various balance constraints in the school zoning problem often represent *competing objectives*. Improving balance along one dimension may require compromising another. The Relaxed ReCom approach struggles to navigate these trade-offs effectively, as it aims to satisfy all constraints simultaneously through the biased spanning tree distribution weights. The performance of the Relaxed ReCom approach is very *sensitive to the choice of weights* assigned to each constraint in the biased

through a narrow path. There are also similar imbalances in the statistics for these zones in terms of the number of schools and the shortage in each zone, which are further provided in Appendix

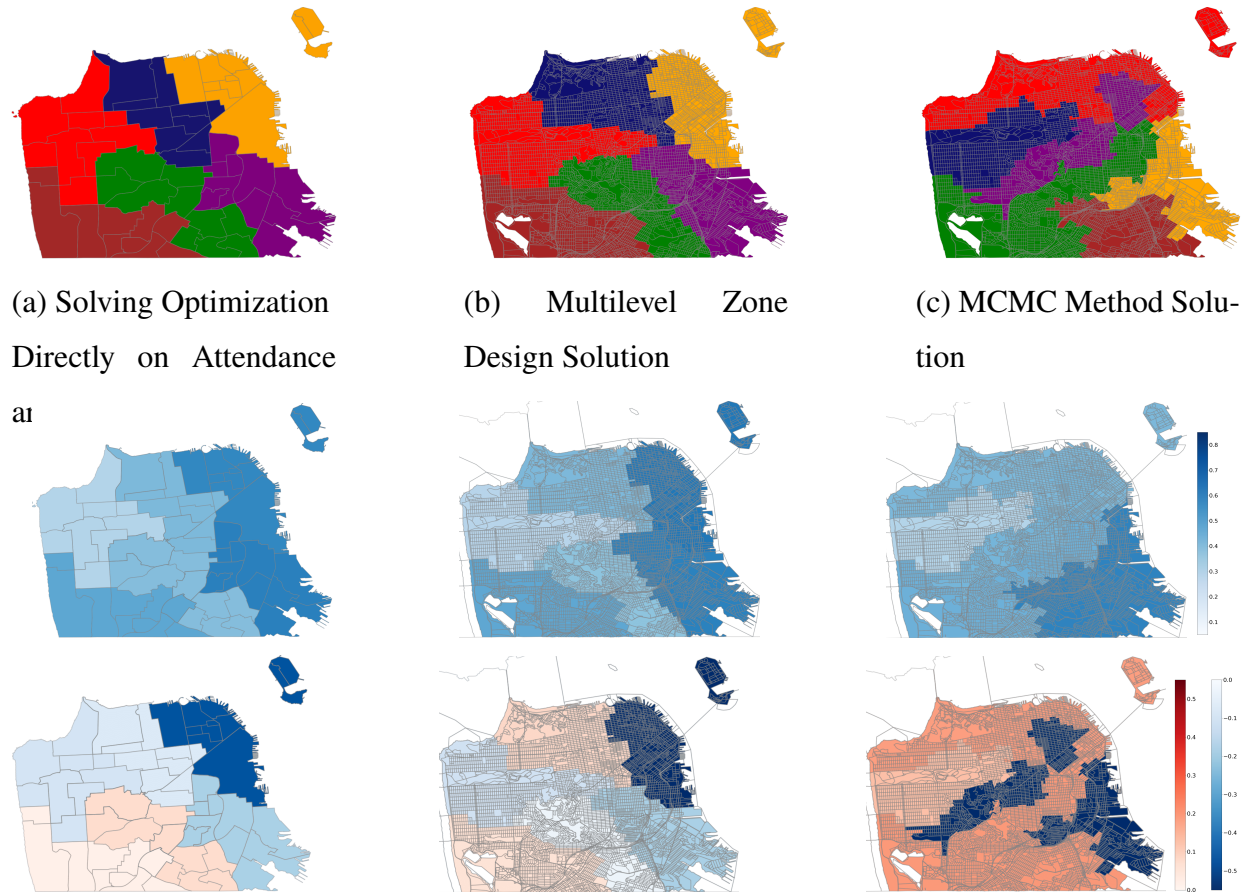


Figure 9: Comparison of 6-zone solutions generated using different methods. The top row shows the zone boundaries, the middle row illustrates the socio-economic balance (percentage of students eligible for free or reduced-price lunch) across zones, and the bottom row depicts the capacity shortage in each zone. In the shortage maps, brighter red indicates higher shortage, while blue represents a seat surplus.

spanning tree distribution. Finding the right set of weights that leads to a well-balanced solution across all dimensions is a non-trivial task and may require extensive tuning and experimentation.

Figure 10 demonstrates that Relaxed ReCom performs relatively well when dealing with a small number of constraints and focusing solely on balancing zone sizes.³⁷ However, the presence of elongated regions, such as in the orange zone, suggests that even with simplified constraints, the MCMC approach may not be the most effective for achieving highly compact and well-structured zones compared to optimization methods.

³⁷ Setting $w_{shortage} = w_{Sch} = w_{FRL} = 0$

ReCom for School Zoning The original ReCom algorithm creates partitions that balance population within a permitted tolerance. For school zoning, we extend this approach to satisfy multiple balance constraints within predefined thresholds. However, the restrictive nature of these constraints significantly reduces the number of feasible cuts at each step. This limitation often leads the algorithm to quickly exhaust valid moves, resulting in local optima traps.




Figure 10: Enforcing only spanning tree balance constraints, using Relaxed ReCom

Simulated Annealing. To address the challenges faced by Relaxed ReCom, we applied simulated annealing to our problem. Starting with power weights of 0 in the biased spanning tree distribution, we sampled for 5,000 steps, then gradually increased the weights to their final values over 10,000 steps, and maintained the final weights for an additional 5,000 steps before taking a sample. However, the solutions obtained showed only marginal improvements in balancing metrics across zones.

5.4. Redrawing Attendance Area Boundaries, and Drawbacks of Multilevel Approach

As part of the school zoning project, SFUSD explored the idea of creating 59 zones, each containing a single school. This concept of neighborhood zones resembles redrawing attendance boundaries, where students are assigned to schools within their designated zones. The objective was to design self-contained zones with adequate capacity to serve the students residing within them while ensuring a balanced student distribution across all zones. To accomplish this, it was necessary to solve the optimization problem directly at the block level.³⁸

We employed the recursive approach in [3.2.3](#) to solve the problem at the block level. The new neighborhood zones (known as ‘attendance area zones’) can be seen in [Figure 11](#). The multilevel approach described in [3.2.2](#) was unsuitable for creating 59 zones, since initializing the process with a zoning solution from either the attendance area or block group level often led to getting stuck at local optima.³⁹

³⁸ The block groups’ relative size compared to the desired smaller 59 zones is quite large, restricting the optimization model’s ability to find an optimal block group arrangement that satisfies the balance and compactness constraints.

³⁹ The optimization model struggled to find a well-balanced and compact zoning solution due to the limited flexibility in adjusting zone boundaries at the lower levels of the hierarchy.

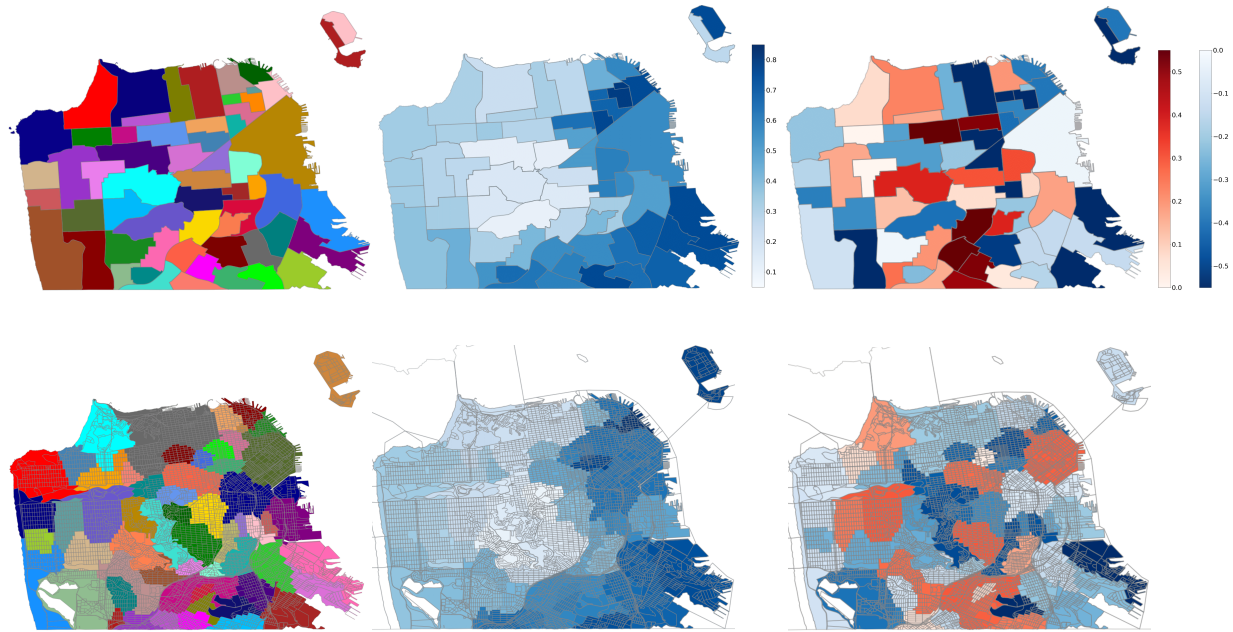


Figure 11: Comparison of the newly generated 59 attendance areas (bottom row) and the historical attendance areas (top row). The left column shows the zone boundaries for each solution. The middle column illustrates the socio-economic balance, depicting the percentage of students eligible for free or reduced-price lunch in each zone. The right column presents the capacity shortage in each zone, with brighter red indicating higher shortage and blue representing a seat surplus.

5.5. Evaluation of Selected Zones

We evaluate the performance of the selected candidate zone maps using an assignment simulation engine, which simulates student preferences over programs within a zone using a multinomial logit utility model trained on historical choice data (92). For the 6, 7, 10, 13, and 18 zone solutions (Appendix Figure 15), we use the same priority and reserve structure from our companion paper on the policy decision, including sibling and equity priorities, as well as reserves based on a composite diversity measure (92). We also restrict students' choices to schools within their assigned zones. For the new attendance areas (59 zones), we run the current policy as a benchmark comparison. The details of this policy can be found in (92).

The resulting assignments are evaluated based on the district's policy priorities of diversity, proximity, and choice metrics to assess the disruption to current choice patterns. Table 1 shows the post-assignment metrics for each zoning solution. We compare the selected assignments to a Status Quo policy, which uses the assignment algorithm that was in place in 2023-24 with student pref-

ferences given by the estimated choice model. Diversity measurements include the percentage of mid-to-high poverty schools (+15% FRL Schools) and the percentage of historically underserved ethnic groups (African American, Latinx, or Pacific Islander) in these schools (AALPI in +15% FRL), motivated by literature suggesting the intersection of these groups in high poverty schools is particularly harmful (103). For proximity, we measure students' average distance to their assigned school (Avg Distance (miles)), and percentages of students assigned within a 'walking distance' of 0.5 miles (Distance \leq 0.5 miles). Choice metrics include the percentage of students receiving one of their top 3 choices (Rank Top 3).

Proximity: All the selected zoning solutions improve the average distance to assigned schools compared to the Status Quo. However, only the 18-zone policy significantly increases the percentage of students assigned to a school within 0.5 miles. Smaller zones generally perform better than larger zones in terms of proximity metrics.

Diversity: Larger zones perform better than smaller zones with respect to diversity. The 6-Zones+Reserves and 7-Zones+Reserves policies are the only ones that decrease both the number of high-poverty schools and the percentage of historically underserved ethnic groups in these schools compared to the Status Quo. Smaller zones perform significantly worse on diversity metrics.

Choice: All zone-based policies restrict choice compared to the current policy, as they limit students' options to schools within their designated zones, which may not include their most preferred programs. However, larger zones allow families to access their preferred school options better than the original attendance area zones.

Updated Attendance Areas: The newly proposed attendance area boundaries (59 zones) show marginal improvements across all metrics compared to the Status Quo, with better performance on proximity and diversity while maintaining the same choice metrics.

6. Conclusion

Designing multi-school zones for student assignment presents significant challenges on both computational and human dimensions. Drawing from the school neighborhood design and political districting literature, we propose a heuristic optimization approach to generate zones for the San Francisco Unified School District. Inspired by the human-computer interaction literature, we develop tools to help learn optimization specifications and support the decision-making process.

Zones pose a considerable design challenge due to the difficulty in conceptualizing their impact, the potential for politically charged design processes, and the computational complexity involved in

		6-Zones + Reserves	7-Zones + Reserves	10-Zones + Reserves	13-Zones + Reserves	18-Zones + Reserves	Updated Attendance Areas	Status- Quo
Proximity	<i>Avg Distance (miles)</i>	1.30	1.29	1.23	1.21	1.21	1.30	1.31
	<i>Distance \leq 0.5 miles</i>	31.7%	33.9%	35.5%	35.5%	36.4%	36%	35.6%
Diversity	<i>+15% FRL Schools</i>	11.2	13.1	14.9	14.6	15.1	12	12.4
	<i>AALPI in +15% FRL</i>	14.9%	17.1%	21%	21.5%	23.8%	17.9%	18.6%
Choice	<i>Rank Top 3</i>	81.3%	81.1%	80.2%	77.2%	75.5%	88.0%	88.0%

Table 1: Average assignment metrics using preferences generated by choice model for counterfactual policies, and the student match of our simulation of the 2023-24 policy.

their development. Despite these challenges, our preliminary simulation results suggest that well-designed zones offer a promising avenue for improving school diversity and student outcomes.

To reach a final community-supported decision, substantial collaboration between researchers and stakeholders is necessary. By leveraging computational techniques and engaging in participatory design processes, we can work towards developing zone-based assignment policies that promote equity and enhance educational opportunities for all students in the San Francisco Unified School District.

References

- [1] Subhodip Biswas, Fanglan Chen, Zhiqian Chen, Chang-Tien Lu, and Naren Ramakrishnan. Incorporating domain knowledge into memetic algorithms for solving spatial optimization problems. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, pages 25–35, 2020.
- [2] Subhodip Biswas. *Spatial Optimization Techniques for School Redistricting*. PhD thesis, Virginia Tech, 2022.
- [3] Nabeel Gillani, Doug Beeferman, Christine Vega-Pourheydarian, Cassandra Overney, Pascal Van Hentenryck, and Deb Roy. Redrawing attendance boundaries to promote racial and ethnic diversity in elementary schools. *arXiv preprint arXiv:2303.07603*, 2023.
- [4] Hernán G Abeledo and Uriel G Rothblum. Courtship and linear programming. *Linear algebra and its applications*, 216:111–124, 1995.
- [5] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [6] Bjørn Nygreen. European assembly constituencies for wales-comparing of methods for solving a political districting problem. *Mathematical Programming*, 42:159–169, 1988.
- [7] Robert S Garfinkel and George L Nemhauser. Optimal political districting by implicit enumeration techniques. *Management Science*, 16(8):B–495, 1970.

- [8] Wes Gurnee and David B Shmoys. Fairmandering: A column generation heuristic for fairness-optimized political districting. In *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*, pages 88–99. SIAM, 2021.
- [9] Anuj Mehrotra, Ellis L Johnson, and George L Nemhauser. An optimization based heuristic for political districting. *Management Science*, 44(8):1100–1114, 1998.
- [10] Daryl DeFord, Moon Duchin, and Justin Solomon. Recombination: A family of markov chains for redistricting. *Harvard Data Science Review*, 3(1):3, 2021.
- [11] Zhanzhan Zhao, Cyrus Hettle, Swati Gupta, Jonathan Christopher Mattingly, Dana Randall, and Gregory Joseph Herschlag. Mathematically quantifying non-responsiveness of the 2021 georgia congressional districting plan. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–11. 2022.
- [12] Douglas M King, Sheldon H Jacobson, and Edward C Sewell. The geo-graph in practice: creating united states congressional districts from census blocks. *Computational Optimization and Applications*, 69:25–49, 2018.
- [13] Felipe Caro, Takeshi Shirabe, Monique Guignard, and Andrés Weintraub. School redistricting: Embedding gis tools with integer programming. *Journal of the Operational Research Society*, 55(8):836–849, 2004.
- [14] Jeroen CJH Aerts, Erwin Eisinger, Gerard BM Heuvelink, and Theodor J Stewart. Using linear integer programming for multi-site land-use allocation. *Geographical analysis*, 35(2):148–169, 2003.
- [15] Hamidreza Validi and Austin Buchanan. Political districting to minimize cut edges. *Mathematical Programming Computation*, 14(4):623–672, 2022.
- [16] Nikhil Garg, Wes Gurnee, David Rothschild, and David Shmoys. Combatting gerrymandering with social choice: The design of multi-member districts. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 560–561, 2022.
- [17] Rahul Swamy, Douglas M King, and Sheldon H Jacobson. Multiobjective optimization for politically fair districting: A scalable multilevel approach. *Operations Research*, 71(2):536–562, 2023.
- [18] Rahul Swamy, Douglas M King, Ian G Ludden, Kiera W Dobbs, and Sheldon H Jacobson. A practical optimization framework for political redistricting: A case study in arizona. *Socio-Economic Planning Sciences*, page 101836, 2024.
- [19] Brian W Kernighan and Shen Lin. An efficient heuristic procedure for partitioning graphs. *The Bell system technical journal*, 49(2):291–307, 1970.
- [20] Bruce Hendrickson, Robert W Leland, et al. A multi-level algorithm for partitioning graphs. *SC*, 95(28):1–14, 1995.
- [21] Aydın Buluç, Henning Meyerhenke, Ilya Safro, Peter Sanders, and Christian Schulz. *Recent advances in graph partitioning*. Springer, 2016.
- [22] Chris Walshaw. A multilevel approach to the travelling salesman problem. *Operations research*, 50(5):862–877, 2002.

-
- [23] Chris Walshaw. A multilevel algorithm for force-directed graph drawing. In *Graph Drawing: 8th International Symposium, GD 2000 Colonial Williamsburg, VA, USA, September 20–23, 2000 Proceedings* 8, pages 171–182. Springer, 2001.
- [24] Harry A Levin and Sorelle A Friedler. Automated congressional redistricting. *Journal of Experimental Algorithms (JEA)*, 24:1–24, 2019.
- [25] Jörg Kalcsics, Stefan Nickel, and Michael Schröder. Towards a unified territorial design approach—applications, algorithms and gis integration. *Top*, 13:1–56, 2005.
- [26] Andrew Spann, Daniel Kane, and Dan Gulotta. Electoral redistricting with moment of inertia and diminishing halves models. *The UMAP Journal*, 28(3):281–299, 2007.
- [27] Sommer Gentry, Eric Chow, Allan Massie, and Dorry Segev. Gerrymandering for justice: redistricting us liver allocation. *Interfaces*, 45(5):462–480, 2015.
- [28] M Angélica Salazar-Aguilar, Roger Z Ríos-Mercado, and José Luis González-Velarde. A bi-objective programming model for designing compact and balanced territories in commercial districting. *Transportation Research Part C: Emerging Technologies*, 19(5):885–895, 2011.
- [29] Douglas M King, Sheldon H Jacobson, Edward C Sewell, and Wendy K Tam Cho. Geo-graphs: an efficient model for enforcing contiguity and hole constraints in planar graph partitioning. *Operations Research*, 60(5):1213–1228, 2012.
- [30] Johannes Ohrlein and Jan-Henrik Haunert. A cutting-plane method for contiguity-constrained spatial aggregation. *Journal of Spatial Information Science*, (15):89–120, 2017.
- [31] H Peyton Young. Measuring the compactness of legislative districts. *Legislative Studies Quarterly*, pages 105–115, 1988.
- [32] Moon Duchin and Bridget Eileen Tenner. Discrete geometry for electoral geography. *arXiv preprint arXiv:1808.05860*, 2018.
- [33] Kiera W Dobbs, Douglas M King, and Sheldon H Jacobson. Redistricting optimization with recombination: A local search case study. *Computers & Operations Research*, 160:106369, 2023.
- [34] Nuno Mota, Negar Mohammadi, Palash Dey, Krishna P Gummadi, and Abhijnan Chakraborty. Fair partitioning of public resources: Redrawing district boundary to minimize spatial inequality in school funding. In *Proceedings of the Web Conference 2021*, pages 646–657, 2021.
- [35] Michael Harwell and Brandon LeBeau. Student eligibility for a free lunch as an ses measure in education research. *Educational researcher*, 39(2):120–131, 2010.
- [36] Lisa M Nicholson, Sandy J Slater, Jamie F Chriqui, and Frank Chaloupka. Validating adolescent socioeconomic status: Comparing school free or reduced price lunch with community measures. *Spatial Demography*, 2:55–65, 2014.

- [37] Thurston Domina, Quentin Brummet, Nikolas Pharris-Ciurej, Sonya R Porter, Andrew Penner, Emily Penner, and Tanya Sanabria. Capturing more than poverty: School free and reduced-price lunch data and household income. *US Census Bureau, Center for Administrative Records Research and Applications. December 2017* <https://www.census.gov/content/dam/Census/library/working-papers/2017/adrm/carra-wp-2017-09.pdf>, 2017.
- [38] Dimitris Bertsimas, Patrick Jaillet, and Sébastien Martin. Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1):143–162, 2019.
- [39] Isidoros Marampoutis, Marina Vinot, and Lorraine Trilling. Multi-objective vehicle routing problem with flexible scheduling for the collection of refillable glass bottles: A case study. *EURO journal on decision processes*, 10:100011, 2022.
- [40] Francesca Rossi, Peter Van Beek, and Toby Walsh. *Handbook of constraint programming*. Elsevier, 2006.
- [41] Louis Perron and Florent Didier. Cpaior 2020 master class: Constraint programming, 2020. <https://youtu.be/lmy1ddn4cyw>.
- [42] Alexander Bockmayr and John N Hooker. Constraint programming. *Handbooks in Operations Research and Management Science*, 12:559–600, 2005.
- [43] John N Hooker and W-J van Hoes. Constraint programming and operations research. *Constraints*, 23:172–195, 2018.
- [44] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. Automated redistricting simulation using markov chain monte carlo. *Journal of Computational and Graphical Statistics*, 29(4):715–728, 2020.
- [45] Jonathan C Mattingly and Christy Vaughn. Redistricting and the will of the people. *arXiv preprint arXiv:1410.8796*, 2014.
- [46] Lucy Chenyun Wu, Jason Xiaotian Dou, Danny Sleator, Alan Frieze, and David Miller. Impartial redistricting: A markov chain approach. *arXiv preprint arXiv:1510.03247*, 2015.
- [47] Gregory Herschlag, Robert Ravier, and Jonathan C Mattingly. Evaluating partisan gerrymandering in wisconsin. *arXiv preprint arXiv:1709.01596*, 2017.
- [48] Benjamin Fifield, Michael Higgins, Kosuke Imai, and Alexander Tarr. A new automated redistricting simulator using markov chain monte carlo. *Work. Pap., Princeton Univ., Princeton, NJ*, 2015.
- [49] Maria Chikina, Alan Frieze, Jonathan C Mattingly, and Wesley Pegden. Separating effect from significance in markov chain tests. *Statistics and Public Policy*, 7(1):101–114, 2020.
- [50] Gregory Herschlag, Han Sung Kang, Justin Luo, Christy Vaughn Graves, Sachet Bangia, Robert Ravier, and Jonathan C Mattingly. Quantifying gerrymandering in north carolina. *Statistics and Public Policy*, 7(1):30–38, 2020.
- [51] Eric A Autry, Daniel Carter, Gregory J Herschlag, Zach Hunter, and Jonathan C Mattingly. Metropolized multiscale forest recombination for redistricting. *Multiscale Modeling & Simulation*, 19(4):1885–1914, 2021.

-
- [52] Gill v. whitford. No. 16-1161, 585 U.S. 1.
- [53] Common cause v. lewis. Supreme Court of North Carolina.
- [54] Rucho v. common cause. 139 S. Ct. 2484.
- [55] League of women voters of pennsylvania v. commonwealth of pennsylvania. Pennsylvania Supreme Court.
- [56] Moses Charikar, Paul Liu, Tianyu Liu, and Thuy-Duong Vuong. On the complexity of sampling redistricting plans. *arXiv preprint arXiv:2206.04883*, 2022.
- [57] Maria Chikina, Alan Frieze, and Wesley Pegden. Assessing significance in a markov chain without mixing. *Proceedings of the National Academy of Sciences*, 114(11):2860–2864, 2017.
- [58] Ariel D Procaccia and Jamie Tucker-Foltz. Compact redistricting plans have many spanning trees. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3754–3771. SIAM, 2022.
- [59] Daryl DeFord and Moon Duchin. Redistricting reform in virginia: Districting criteria in context. *Virginia Policy Review*, 12(2):120–146, 2019.
- [60] Hugo A Akitaya, Andrei Gonczi, Diane L Souvaine, Csaba D Tóth, and Thomas Weighill. Reconfiguration of polygonal subdivisions via recombination. *arXiv preprint arXiv:2307.00704*, 2023.
- [61] Stuart Wayland. Quantifying gerrymandering with simulated annealing. *arXiv preprint arXiv:2209.00624*, 2022.
- [62] Sachet Bangia, Christy Vaughn Graves, Gregory Herschlag, Han Sung Kang, Justin Luo, Jonathan C Mattingly, and Robert Ravier. Redistricting: Drawing the line. *arXiv preprint arXiv:1704.03360*, 2017.
- [63] Vincent Cohen-Addad, Philip N Klein, and Dániel Marx. On the computational tractability of a geographic clustering problem arising in redistricting. *arXiv preprint arXiv:2009.00188*, 2020.
- [64] Eric A Autry, Daniel Carter, Gregory Herschlag, Zach Hunter, and Jonathan C Mattingly. Multi-scale merge-split markov chain monte carlo for redistricting.(2020). URL: <https://arxiv.org/abs>, 2008.
- [65] Marija Vucelja. Lifting—a nonreversible markov chain monte carlo algorithm. *American Journal of Physics*, 84(12):958–968, 2016.
- [66] Manon Michel. *Irreversible Markov chains by the factorized Metropolis filter: algorithms and applications in particle systems and spin models*. PhD thesis, Université Paris sciences et lettres, 2016.
- [67] Gregory Herschlag, Jonathan C Mattingly, Matthias Sachs, and Evan Wyse. Non-reversible markov chain monte carlo for sampling of districting maps. *arXiv preprint arXiv:2008.07843*, 2020.
- [68] Cory McCartan and Kosuke Imai. Sequential monte carlo for sampling balanced and compact redistricting plans. *The Annals of Applied Statistics*, 17(4):3300–3323, 2023.
- [69] Maria CV Nascimento and Andre CPLF De Carvalho. Spectral methods for graph clustering—a survey. *European Journal of Operational Research*, 211(2):221–231, 2011.

- [70] Satu Elisa Schaeffer. Survey: Graph clustering. *comput. sci. rev.* 1, 1 (aug. 2007), 27–64, 2007.
- [71] David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 296–303, 1996.
- [72] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, Cynthia Vinzant, and Thuy-Duong Vuong. Log-concave polynomials iv: approximate exchange, tight mixing times, and near-optimal sampling of forests. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 408–420, 2021.
- [73] Karen Levy, Kyla E Chasalow, and Sarah Riley. Algorithms and decision-making in the public sector. *Annual Review of Law and Social Science*, 17:309–334, 2021.
- [74] Arthur Delarue, Zhen Lian, and Sebastien Martin. Algorithmic precision and human decision: A study of interactive optimization for school schedules. *Available at SSRN 4324076*, 2023.
- [75] Ningyuan Chen, Ming Hu, and Wenhao Li. Algorithmic decision-making safeguarded by human knowledge. *arXiv preprint arXiv:2211.11028*, 2022.
- [76] Andrew M Davis, Shawn Mankad, Charles J Corbett, and Elena Katok. The best of both worlds: Machine learning and behavioral science in operations management. *Available at SSRN 4258273*, 2022.
- [77] Michael J Muller and Sarah Kuhn. Participatory design. *Communications of the ACM*, 36(6):24–28, 1993.
- [78] Elizabeth B-N Sanders and Pieter Jan Stappers. Probes, toolkits and prototypes: three approaches to making in codesigning. *CoDesign*, 10(1):5–14, 2014.
- [79] Jesper Simonsen and Toni Robertson. *Routledge international handbook of participatory design*, volume 711. Routledge New York, 2013.
- [80] Judith Gregory. Scandinavian approaches to participatory design. *International Journal of Engineering Education*, 19(1):62–74, 2003.
- [81] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–23, 2023.
- [82] Gene Rowe and Lynn J Frewer. Public participation methods: a framework for evaluation. *Science, technology, & human values*, 25(1):3–29, 2000.
- [83] Yves Cabannes. The impact of participatory budgeting on basic services: municipal practices and evidence from the field. *Environment and Urbanization*, 27(1):257–284, 2015.
- [84] Rachel Swaner. Trust matters: Enhancing government legitimacy through participatory budgeting. *New Political Science*, 39(1):95–108, 2017.
- [85] Julia Abelson, John Eyles, Christopher B McLeod, Patricia Collins, Colin McMullan, and Pierre-Gerlier Forest. Does deliberation make a difference? results from a citizens panel study of health goals priority setting. *Health policy*, 66(1):95–106, 2003.

-
- [86] Evelien AG Joosten, Laura DeFuentes-Merillas, GH De Weert, Tom Sensky, CPF Van Der Staak, and Cor AJ de Jong. Systematic review of the effects of shared decision-making on patient satisfaction, treatment adherence and health status. *Psychotherapy and psychosomatics*, 77(4):219–226, 2008.
- [87] John Elliot. *Action research for educational change*. McGraw-Hill Education (UK), 1991.
- [88] Gillian R Hayes. Knowing by doing: action research as an approach to hci. In *Ways of Knowing in HCI*, pages 49–68. Springer, 2014.
- [89] Lauge Baungaard Rasmussen. Action research—scandinavian experiences. *AI & SOCIETY*, 18:21–43, 2004.
- [90] Kim M Unertl, Chris L Schaefbauer, Terrance R Campbell, Charles Senteio, Katie A Siek, Suzanne Bakken, and Tiffany C Veinot. Integrating community-based participatory research and informatics approaches to improve the engagement and health of underserved populations. *Journal of the American Medical Informatics Association*, 23(1):60–73, 2016.
- [91] Nina B Wallerstein and Bonnie Duran. Using community-based participatory research to address health disparities. *Health promotion practice*, 7(3):312–323, 2006.
- [92] Maxwell Allman, Itai Ashlagi, Irene Lo, Juliette Love, Katherine Mentzer, Lulabel Ruiz-Setz, and Henry O’Connell. Designing school choice for diversity in the san francisco unified school district. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 290–291, 2022.
- [93] John R Hauser and Birger Wernerfelt. An evaluation cost model of consideration sets. *Journal of consumer research*, 16(4):393–408, 1990.
- [94] Mahsa Derakhshan, Negin Golrezaei, Vahideh Manshadi, and Vahab Mirrokni. Product ranking on online platforms. *Management Science*, 68(6):4024–4041, 2022.
- [95] Ali Aouad and Danny Segev. Display optimization for vertically differentiated locations under multinomial logit preferences. *Management Science*, 67(6):3519–3550, 2021.
- [96] Guillermo Gallego, Anran Li, Van-Anh Truong, and Xinshang Wang. Approximation algorithms for product framing and pricing. *Operations Research*, 68(1):134–160, 2020.
- [97] Ruxian Wang and Ozge Sahin. The impact of consumer search cost on assortment planning and pricing. *Management Science*, 64(8):3649–3666, 2018.
- [98] Charles Hodgson and Gregory Lewis. You can lead a horse to water: Spatial learning and path dependence in consumer search. Technical report, National Bureau of Economic Research, 2023.
- [99] Pedro M Gardete and Megan Hunter. Avoiding lemons in search of peaches: Designing information provision. Technical report, Stanford University, Graduate School of Business, 2018.
- [100] Malika Korganbekova and Cole Zuber. Balancing user privacy and personalization. *Work in progress*, 2023.
- [101] Pranav Jindal and Anocha Aribarg. The importance of price beliefs in consumer search. *Journal of Marketing Research*, 58(2):321–342, 2021.

- [102] Patricia Gayá Wicks, Peter Reason, and Hilary Bradbury. Living inquiry: Personal, political and philosophical groundings for action research practice. *The Sage handbook of action research: Participative inquiry and practice*, 2:15–30, 2008.
- [103] Sean F Reardon, Ericka Weathers, Erin Fahle, Heewon Jang, and Demetra Kalogrides. Is separate still unequal? new evidence on school segregation and racial academic achievement gaps. 2019.
- [104] Peng Shi. Optimal priority-based allocation mechanisms. *Management Science*, 68(1):171–188, 2022.

7. Computational Methods

7.1. Arbitrary Centroids

In this subsection, we provide an integer program in which centroids are not predefined anymore, but instead centroids are variables that could be any unit, while satisfying all the other balance/contiguity/etc constraints. This will allow our integer program to find a larger class of solutions, and maybe find zonings with given constraints in which we previously deemed to restraining to have any solution for. Unfortunately this integer program has substantially more variables, which makes it intractable to solve.

Let y_z be a binary decision variable indicating whether unit z is the centroid of a zone or not, $\forall z \in S$. Let $x_{u,z}$ be the decision variable indicating whether unit u is part of a zone with unit z as its centroid-unit or not, $\forall u, z \in U$. In this new integer program, we enforce all the previously introduced constraints, but anywhere among constraints [1](#) - [8](#) where we had the term “ $z \in Z$ ” we now replace it with “ $z \in U$ ”. We additionally enforce the following constraints.

$$\sum_{z \in U} y_z = |Z| \quad (12)$$

Constraint [12](#) makes sure we have exactly $|Z|$ centroids.

$$x_{u,z} \leq y_z \quad \forall u, z \in U \quad (13)$$

Constraint [13](#) makes sure we only assign a unit u to a zone with centroid z if unit z is considered as a centroid.

7.2. Combining Zone Optimization with Choice

In order to combine zone optimization and choice into a single, tractable math program, we use a priority optimization formulation similar to Shi ([104](#)).

Let S be the set of schools, and let T be the set of student types, where students of type $t \in T$ have utilities drawn from distribution F_t . Let $n_{u,t}$ be the number of students in unit u of type t $\forall u \in U, t \in T$, and let $\text{Util}_t(S') := \mathbb{E}_{f \sim F_t} [\max_{u \in S} f_u]$ be the expected utility of a type t student with budget set $S' \subseteq S$. We further define $P_t(s, S') := 1\{s \in S'\} \mathbb{P}_{f \sim F_t} (f_s = \max_{s' \in S'} f_{s'})$ is the probability a student of type t 's favorite school in S' is s . Let variable $y_{t,u,S'}$ be the probability that a student of type t living in unit u has budget set S' .

$$\max_y \sum_{u,t,S'} n_{u,t} \text{Util}_t(S') y_{t,u,S'} \quad (14)$$

$$\text{s.t.} \quad y_{t,u,S'} \geq 0 \quad \forall u \in U, t \in T, S' \subseteq S \quad (15)$$

$$\text{(Capacity)} \quad \sum_{t,u,S'} n_{u,t} P_t(s, S') y_{t,u,S'} \leq c_s \quad \forall u \in U, t \in T, s \in S \quad (16)$$

$$\text{(Valid probabilities)} \quad \sum_{S' \subseteq S} y_{t,u,S'} = 1 \quad \forall u \in U, t \in T \quad (17)$$

7.3. Combining zone optimization and priority optimization and choice.

Suppose we want to combine zones with priority optimization and choice by directly using the above math programs. Then we need to introduce additional constraints on valid probabilities $y_{t,u,S'}$. Specifically, we need to ensure that a student's budget set is contained in their zone, i.e.

$$y_{t,u,S'} > 0 \text{ only if } S' \subseteq \{\text{schools in } u\text{'s zone}\}.$$

A school s is in student u 's zone if and only if $\sum_{z \in Z} x_{s,z} x_{u,z} = 1$, and otherwise $\sum_{z \in Z} x_{s,z} x_{u,z} = 0$. So we can ensure each students' budget set is contained in their zone using the *quadratic* constraints

$$y_{t,u,S'} \leq \sum_{z \in Z} x_{s,z} x_{u,z} \quad \forall t \in T, S' \subseteq S, s \in S. \quad (18)$$

Making the quadratic constraints linear. In this subsection, we provide an alternative formulation of the zone optimization program that increases the number of variables and constraints (by a factor of $|Z|/|S|$) but allows us to characterize valid probabilities using a linear constraint. Let $x_{u,s}$ be a decision variable indicating whether unit u is part of the same zone as school s , defined for all $u \in U, s \in S$. We write the following optimization program.

$$\min \sum_{u,t,S'} n_{u,t} F_t(S') y_{t,u,S'} + \sum_{u \in U, v \in N(u)} b_{u,v} \quad (19)$$

$$\text{s.t.} \quad |x_{u,s} - x_{v,s}| \leq 2 - (x_{u,s'} + x_{v,s'}) \quad \forall v, u \in U, s, s' \in S \quad (20)$$

$$y_{t,u,S'} \leq x_{u,s} \quad \forall u \in U, t \in T, S' \subseteq S, s \in S' \quad (21)$$

$$y_{t,u,S'} \geq 0 \quad \forall u \in U, t \in T, S' \subseteq S \quad (22)$$

$$\sum_{t,u,S'} n_{u,t} P_t(s, S') y_{t,u,S'} \leq c_s \quad \forall u \in U, t \in T, s \in S \quad (23)$$

$$\sum_{S' \subseteq S} y_{t,u,S'} = 1 \quad \forall u \in U, t \in T. \quad (24)$$

$$x_{u,s} \leq \sum_{v: v \in N(u), d_{v,s} \leq d_{u,s}} x_{v,s} \quad \forall u \in U, s \in S \quad (25)$$

$$0.15 \cdot \sum_{u \in U} n_u x_{u,s} \geq \left| \sum_{u \in U} (n_u - q_u) x_{u,s} \right| \quad \forall s \in S \quad (26)$$

$$1 \geq \left| \sum_{u \in U} \frac{sch_u}{|Z|} - \sum_{u \in U} sch_u x_{u,s} \right| \quad \forall s \in S \quad (27)$$

$$\sum_{u \in U} f_u x_{u,s} \geq \left(\frac{F}{N} - 0.15 \right) \cdot \sum_{u \in U} n_u x_{u,s} \quad \forall s \in S \quad (28)$$

$$\sum_{u \in U} R_u^k x_{u,s} \geq \left(\frac{R^k}{N} - 0.15 \right) \cdot \sum_{u \in U} n_u x_{u,z} \quad \forall s \in S, k \in K \quad (29)$$

Note that this program is almost identical to combining the previous two math programs and replacing all instances of $x_{u,z}$ with $x_{u,s}$. There are a few notable differences, which are as follows. In the objective, we have replaced the compactness objective with a choice objective (total utility). Constraint [20](#) is the zone constraint. It ensures that if two units u, v are in the same zone as some school s' , then the set of schools in u 's zone is the same as the set of schools in v 's zone (i.e. $x_{u,s} = x_{v,s}$). Otherwise, the units do not share any schools. Constraint [21](#) is the constraint for valid probabilities. It ensures that a student's budget set is contained in their zone. Constraints [22](#) - [24](#) are as in the math program [14](#) in subsection [7.2](#), and all other constraints [25](#)-[29](#) are as in the math program in subsection [3.2](#), with $x_{u,z}$ replaced with $x_{u,s}$; for all $s \in S$ rather than for all $z \in Z$.




Figure 12: Visualizing the Limitations of the Contiguity Constraint in (2)

7.4. Contiguity Constraint

Figure 12 illustrates a case where (2) may fail: unit a does not have any neighbors closer to centroid z than itself. To address this issue, we enforce constraint (2) only for units u and zones z where u has at least one neighbor closer to z .

7.5. Local Search

Starting with the solution obtained from the multilevel process 3.2.2, we apply the Local Search algorithm to further improve the zoning results. The local search process is repeated for a maximum of 5000 iterations or until convergence is reached (whichever occurs first). It is important to note that the Local Search algorithm is designed to operate at the census block level for the geographical units. This is possible because at each iteration, the algorithm only needs to solve a bipartition problem on a subset of the district.

Algorithm 4 Local Search

Input: Current school zoning, P . Let $P(u)$ be the zone for unit u in zoning P .

Output: New school zoning, Q .

- Randomly select two adjacent zones Z_1 and Z_2 in P .
- Merge zones Z_1 and Z_2 to create a single partition U .
- Split the partition U into two zones, \hat{Z}_1 and \hat{Z}_2 .
- Define $Q(u) = \begin{cases} \hat{Z}_1 & \text{if } u \in \hat{Z}_1 \\ \hat{Z}_2 & \text{if } u \in \hat{Z}_2 \\ P(u) & \text{otherwise} \end{cases}$

Return Q

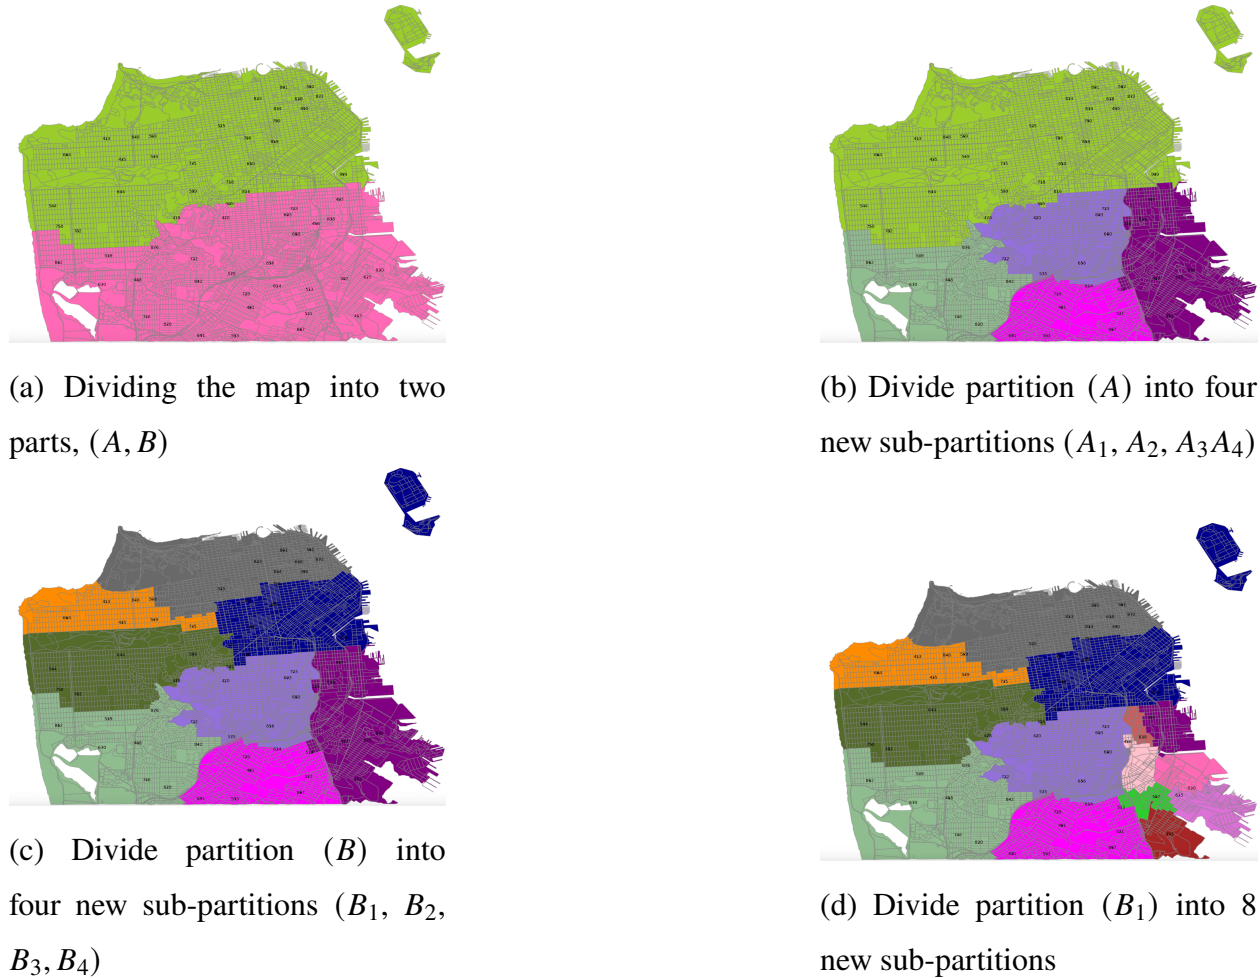


Figure 13: Iterations of the Divide and Conquer process

7.6. Divide-and-Conquer Zoning

We present a recursive algorithm that divides the district into non-overlapping subsections and solves each subproblem separately. The algorithm steps are: (I) Divide the district into two balanced partitions, A and B , minimizing the shortage of school seats and boundary length (Figure 13a). (II) Further divide each partition into four sub-partitions (A_1, A_2, A_3, A_4) and (B_1, B_2, B_3, B_4) ; Figures 13c and 13c), ensuring an equal number of schools per zones, balanced shortage, and minimal boundaries for compactness. (III) Divide each sub-partition $(A_i$ or $B_i)$ into roughly 8 zones⁴⁰, prioritizing zone balance and compactness. (IV) Perform post-processing to improve boundaries.

One limitation of the Divide and Conquer approach is that the zone corners or edges become fixed due to previous partitioning steps when sub-partitions $(A_i$ or $B_i)$ are divided into smaller zones. This

⁴⁰We experimented with different splitting patterns (e.g., 2/4/8, 4/2/8) but focused on the 2/4/8 split for illustration.

restricts flexibility in adjusting zone boundaries and limits solution space variation (Figure 13d). As a result, this approach struggles to achieve a good balance across zones in terms of shortage and socio-economic diversity. Solutions obtained using Divide and Conquer have significantly worse balance compared to the Overlapping Recursive Zoning introduced in section 3.2.3.⁴¹

A solution that appears compact and well-balanced when dividing the district into two partitions may not be optimal when further subdividing partitions into smaller zones. Initial partitioning decisions significantly impact the quality of the final zoning solution. Ideally, a large number of zones should have a circular or compact shape to minimize total boundary length and promote neighborhood cohesion. However, the initial two-partition division may introduce sharp corners or elongated shapes that are not optimal for the final zoning configuration, limiting the ability to achieve desired compactness in subsequent steps.

7.7. Jointly Optimizing Resource Allocation and Zone Design

In this section, we propose a solution to optimize the allocation of educational programs across schools in San Francisco while maintaining the total number of program offerings. The current fixed locations of schools pose a challenge, as the existing distribution of schools and their capacities may not optimally serve the student population. Altering these locations would be prohibitively expensive due to the high costs associated with closing and opening schools. However, while the physical school locations cannot be easily changed, we suggest that modifying the allocation of classrooms within each school can better meet the needs of students.

In SFUSD, kindergarten students apply to 30 special education programs (8 different types) that are crucial for supporting students with disabilities. In the 2019-20 academic year, over 3,400 children in grades K-5 were enrolled in these programs, with 78% identified for Speech or Language Impairment, Specific Learning Disability, Autism, and Other Health Impairment. Due to the wide variety and limited availability of special education programs, the proposed zoning system allows students to be eligible for all such programs district-wide, while restricting their eligibility for GE programs to their assigned zone.

Despite the importance of these programs, several equity gaps exist. Students enrolling in special education programs are more likely to come from marginalized communities. Suboptimal program placement has led to increased travel times and transportation costs for students, and enrollment

⁴¹ Zones may have a shortage of up to 45%, much higher than the 28% shortage achieved by solutions in section 3.2.3.

patterns reveal disparities in access to different Special Education programs based on disability type and race.

SFUSD is obligated to provide transportation for students enrolling in special education programs due to their specific needs and the limited number of programs available across the district. The cost of busing for special education is substantial, amounting to \$26 million per year, primarily because students do not always live close to the schools offering the programs they require. However, the demand for special education programs is not entirely unpredictable, as specific neighborhoods have historically shown consistent demand for particular programs over time.

Proposed Solution. To address these challenges, we propose reallocating programs across schools while maintaining the total number of offerings for each program. By offering programs at schools located close to areas with high demand, we aim to reduce transportation costs. The new problem simultaneously tackles two goals:

- Find a relocation of special education programs that minimizes the total cost of transportation for students enrolling in these programs.
- Divide the district into a specified number of zones, limiting general education (GE) students' choices to schools within their assigned zone.

Incorporating post-choice outcomes, such as transportation costs for special education students, directly into the objective function is not feasible (as shown in Appendix 7.2). Instead, we model the problem as a centralized system that assigns students to special education programs while minimizing total transportation cost. This heuristic provides a good estimate of transportation costs, guiding the effective relocation of programs.

We aim to develop a single optimization model that addresses both Goal 1 (relocating classes) and Goal 2 (finding the best zoning). A combined approach can lead to better solutions that might be missed in a two-step optimization heuristic. By simultaneously relocating programs and finding the best assignment of students to schools for special education, we can achieve a better results.

Let T be the set of special education programs, including Mild/Moderate Autism Focus (AF), Emotionally Disturbed (ED), Spanish Immersion (SE), etc. Let $\text{Distance}_{u,s}$ be the transit distance between unit u and school s . Let Population_u^t be the number of students in unit u that would like enroll in program t . Let $A_{u,s}^t$ be a binary variable indicating whether students of Population_u^t are assigned to school s . To generate zones for GE students while optimizing the transit distance for students assigned to special education programs, we modify the mixed integer program from subsection 3.2.

$$Obj_{SpEd} = \sum_{t \in T} \sum_{u \in U} \sum_{s \in S} \text{Distance}_{u,s} \cdot \text{Population}_u^t \cdot A_{u,s}^t$$

Objective The objective function computes the total transit distance of each special education program student to their matched school, which will be combined with the previous objective value [3](#) for generating zones for GE students. This way we simultaneously optimize the transit distance of students assigned to relocated special education programs classes while finding compact zones for GE students.

$$\sum_{s \in S} A_{u,s}^t = 1 \quad \forall t \in T, u \in U \quad (30)$$

Assignment Feasibility Constraint [30](#) ensures that students of all special education programs, from all neighborhoods, are assigned to a school that offers that program.

In San Francisco, each school can have a maximum of only one class for each special education program. This ensures that the limited number of classes for each program is distributed across different schools rather than concentrated in a single school. In contrast, schools are allowed to have multiple classrooms for general education (GE). To model this specific criteria, we introduce the decision variable q_s^t , which indicates whether school s has a class for students of program t . Let Capacity^t be the capacity of a classroom for program $t \in T$, considering the required equipment and the number of students a teacher can handle for each specialty need.

$$\sum_{u \in U} \text{Population}_u^t \cdot A_{u,s}^t \leq \text{Capacity}^t \cdot q_s^t \quad \forall t \in T, s \in S$$

Assignment Capacity This constraint ensures that for each school s , there are enough seats for all the students with each special education type matched to school s .

Modeling the classrooms for GE is slightly more complex since schools can have multiple GE classrooms. We introduce a separate variable for each GE class in each school. Let $y_{s,z}^i$ be a binary variable indicating whether the i th classroom of school s is offering GE program, and is part of zone z . Let Rooms_s be the number of classrooms in school s .

$$y_{s,z}^i \geq y_{s,z}^{i+1} \quad \forall s \in S, z \in Z, i < \text{Rooms}_s \quad (31)$$

School Classroom Unification Constraint [31](#) ensures that the GE classes in each school are assigned to only one zone, preventing the optimization model from assigning different GE classrooms of a single school to separate zones.

$$x_{s,z} \geq y_{s,z}^1 \quad \forall s \in S, z \in Z \quad (32)$$

Zone Alignment Constraint [32](#) ensures that classes of school s are assigned to zone z only if the area unit in which school s is located is also assigned to the same zone. In the zone optimization model in subsection [3.2](#), we assigned each unit to a zone using decision variables $x_{u,z}$. Now, our GE classrooms within a school must follow the same logic: a school with its GE classes can be assigned to a zone only if that unit area is assigned to that zone. It is sufficient to form the inequality constraint for $y_{s,z}^1$, given the constraint in [31](#).

Note: The introduction of the variables $y_{s,z}^i$ might seem unnecessarily complicated, but they are crucial for finding the zoning solution simultaneously using a linear model. If we were to use a simpler variable q_s^{GE} to represent the number of GE classes in a school s , similar to the approach used for special education classes, we would encounter a quadratic constraint when computing the shortage of each zone. This constraint would limit the value of $\sum_{u \in U} x_{u,z} (q_u^{GE} - \text{Population}_u^{GE})$, which is a product of two decision variables. Quadratic constraints are much more computationally expensive to solve, even though the model would have fewer variables compared to our linear formulation using $y_{s,z}^i$.

$$\text{Rooms}_s = \sum_{t \in T} q_s^t + \sum_{z \in Z} \sum_i y_{s,z}^i \quad \forall s \in S \quad (33)$$

Unchanging Classroom Quota Constraint [33](#) ensures that the total number of classes in each school remains the same, as the physical limits on the number of rooms in the building prevent increasing or decreasing the total number of classes within a school.

Let Classes^t be the total number of classes offered for program t across the district.

$$\text{Classes}^t = \sum_{s \in S} q_s^t \quad \forall t \in T \quad (34)$$

Unchanging Special Education Quota Constraint [34](#) ensures that the total number of classes for each special education program remains constant across the district. At this phase of the work, we do not want to increase or decrease the overall number of programs offered in the district.

$$0.15 \cdot \sum_{u \in U} x_{u,z} \cdot \text{Population}_u^{GE} \geq \left| \sum_{u \in U} \text{Population}_u^{GE} - \text{Capacity}^{GE} \cdot \sum_{z \in Z} \sum_i y_{s,z}^i \right| \quad \forall z \in Z$$

Shortage To generate zones for GE students while optimizing the transit distance for students enrolling in special education programs, we modify the mixed integer program from subsection 3.2. All constraints remain the same as in subsection 3.2, with the exception of the shortage constraint 6, which is redefined as follows. Constraint 7.7 ensures that for each zone z , there are enough seats for GE students assigned to zone z among all the GE classes assigned to that zone:

Finally, we have the following constraints on the decision variables:

$$y_{s,z}^i, q_s^t, A_{u,s}^t \in \{0, 1\} \quad \forall u \in U, z \in Z, s \in S, i < \text{Rooms}_s$$

By incorporating these constraints and objectives into a single optimization model, we can effectively reallocate special education programs across schools in San Francisco while simultaneously drawing optimal zones for general education students. This approach aims to reduce transportation costs for special education programs and improve the overall efficiency of the school system.

Results We analyze the change in the distribution of the Autism Focus (AF) program, based on our optimization model.

Figure 14a illustrates the current location of AF programs in the San Francisco Unified School District. To optimize the distribution of these programs, we run our model with the objective of reducing transit times for AF students while simultaneously reallocating school classes and drawing zone boundaries to divide the district into 6 zones (used here as a sample example). The results of this optimization are presented in Figure 14b.

To further demonstrate the effectiveness of our approach, we analyze the distribution of students who participated in AF programs from 2014–2023. The optimized locations of AF programs exhibit a higher proximity to areas with a high concentration of AF participant students. Moreover, the new AF program locations are not only closer to AF students but also better scattered throughout the district, providing more even coverage. These improvements in proximity and distribution contribute to reduced transit times for AF students, ultimately enhancing the accessibility and efficiency of the special education program.

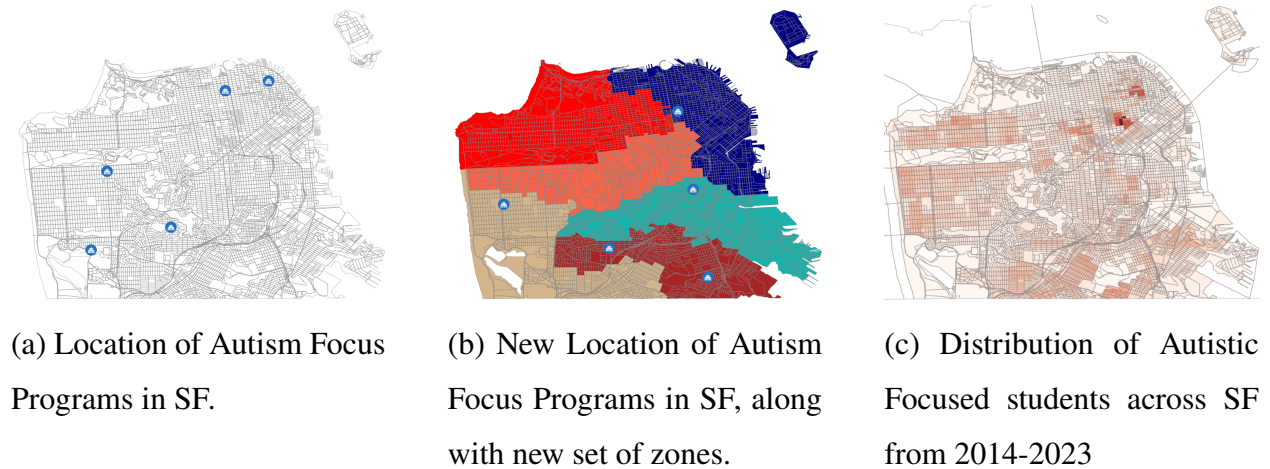


Figure 14

8. Appendix: Results

8.1. Score-Based Filtering Approach

The scoring and filtering process relied on 10 working definitions: 5 diversity goals, 3 proximity goals, and 2 predictability goals.

1. Diversity – Free or reduced price lunch dissimilarity
2. Diversity – Racial dissimilarity
3. Diversity – English language proficiency dissimilarity
4. Diversity – Special education dissimilarity
5. Diversity – Academic performance dissimilarity via level 1 (lowest) standardized test scores
6. Proximity – Average walk time to all schools in zone
7. Proximity – Percent of students within 20 minute walking distance of 2 or more schools
8. Proximity – Percent of students within 20 minute transit distance of 2 or more schools
9. Predictability – General education seat deficit
10. Predictability – Total seat deficit

Each working definition was folded into a total map score via a weighted average. Each diversity objective was 10% of the total score (diversity giving 50% contribution), each proximity score was worth 6.67% of the total score (proximity giving a 20% contribution), and each of the the predictability goals was worth 15% of the total score (predictability contributing 30% of the score).

While effective at narrowing the field of candidate maps, the scoring and filtering method had several shortcomings. First, it does not include a wide variety of stakeholders in the decision-making

process and is not a participatory design approach. The method relies fairly heavily on the judgment of a small group of stakeholders to determine reasonable metrics, weightings, and cutoffs, and to an outsider, these decisions may seem arbitrary. Second, this filtering approach may not capture the full breadth of potential solutions in its results. The score-based filtering stages implicitly impose assumptions about the trade-offs between objectives; specifically, the method preserves solutions in the middle part of the Pareto frontier of diversity, proximity, and predictability. For example, a map that performs well on diversity and proximity might be eliminated for falling in the bottom half of the predictability score. However, eliminating this map may only be a good idea if relative performance is a good marker of viability. If an absolute threshold is a better marker, reasonable maps could be eliminated – extending the previous example, if the map that performs well on diversity and proximity fell into a predictability range that could be addressed by minor capacity planning adjustments, stakeholders may prefer to keep this map in contention.

In response to these concerns, the future filtering process will likely differ markedly from this scoring approach. Current work is underway to find a more participatory and balanced approach to zone filtering.

8.2. Finalized Zones

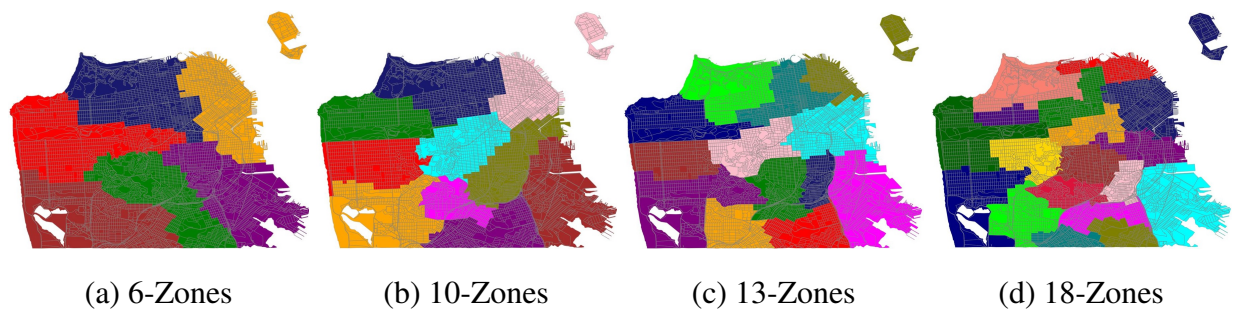


Figure 15: Optimized school assignment zones for 6, 10, 13, and 18 zones in San Francisco.

8.3. Interactive Selection Tool

The interactive selection tool was developed using Google Sheets, a familiar and easy-to-use platform that allows for building interactive web interfaces without extensive web development experience. The choice of platform was inspired by the success of a similar tool used by SFUSD decision-makers to explore trade-offs in school start times (74).

3. Diversity Exploration: How diverse can we make zones?

Building off of our previous results, we can now start to think about diversity. In this tab, you will pick your favorite zones based on diversity, subject to the capacity constraint you set on the previous tab. You can go back and loosen your capacity constraint if you run out of zones here!

Here, you can filter based on the largest allowable fraction of a zone that is made up of one ethnic group (Max Ethnic Group Fraction of Zone) and the maximum FRL eligibility of any zone in the map (Max Zone FRL). Try maximum ethnic group fractions of 35%, 38%, and 40% to get started, and maximum zone FRL of 60%, 63%, and 65%.

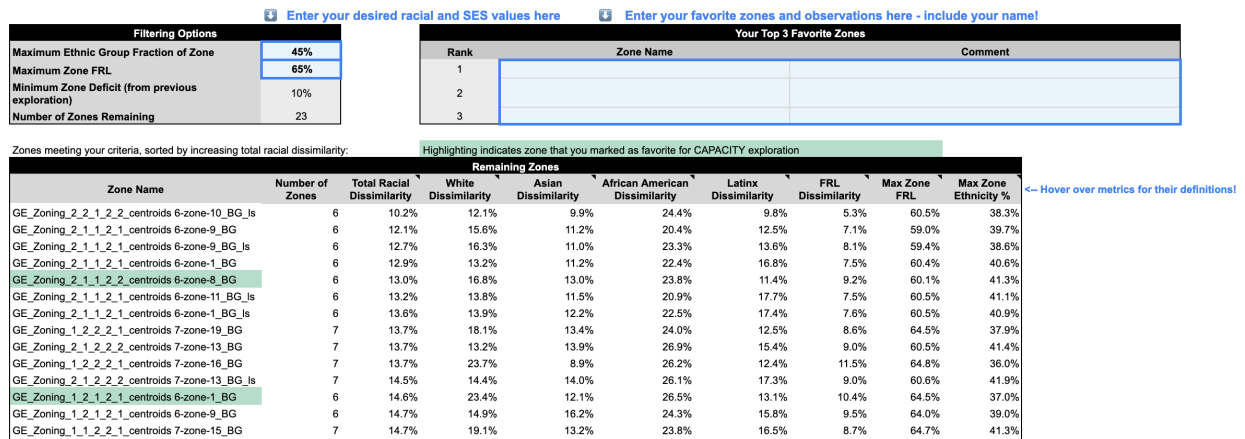


Figure 16: The diversity selection stage.

The tool incorporates several design features to enhance user experience, such as color coding, tooltips, expandable sections for detailed metrics, and a plot comparing the current zone’s performance to other zones in the solution family (Figure 17). Inspired by gerrymandering identification strategies (?), the tool highlights the selected map’s position in a histogram of metrics like racial dissimilarity, illustrating the trade-offs made relative to other maps in the collection. These design choices aim to simplify the complex task of comparing maps.

Initial testing with district staff revealed that the tool was overly complicated and ineffective due to participants’ insufficient familiarity with the problem and the nuances of zone design. Users found the amount of new information overwhelming, and the researchers provided insufficient problem setup and too much information.

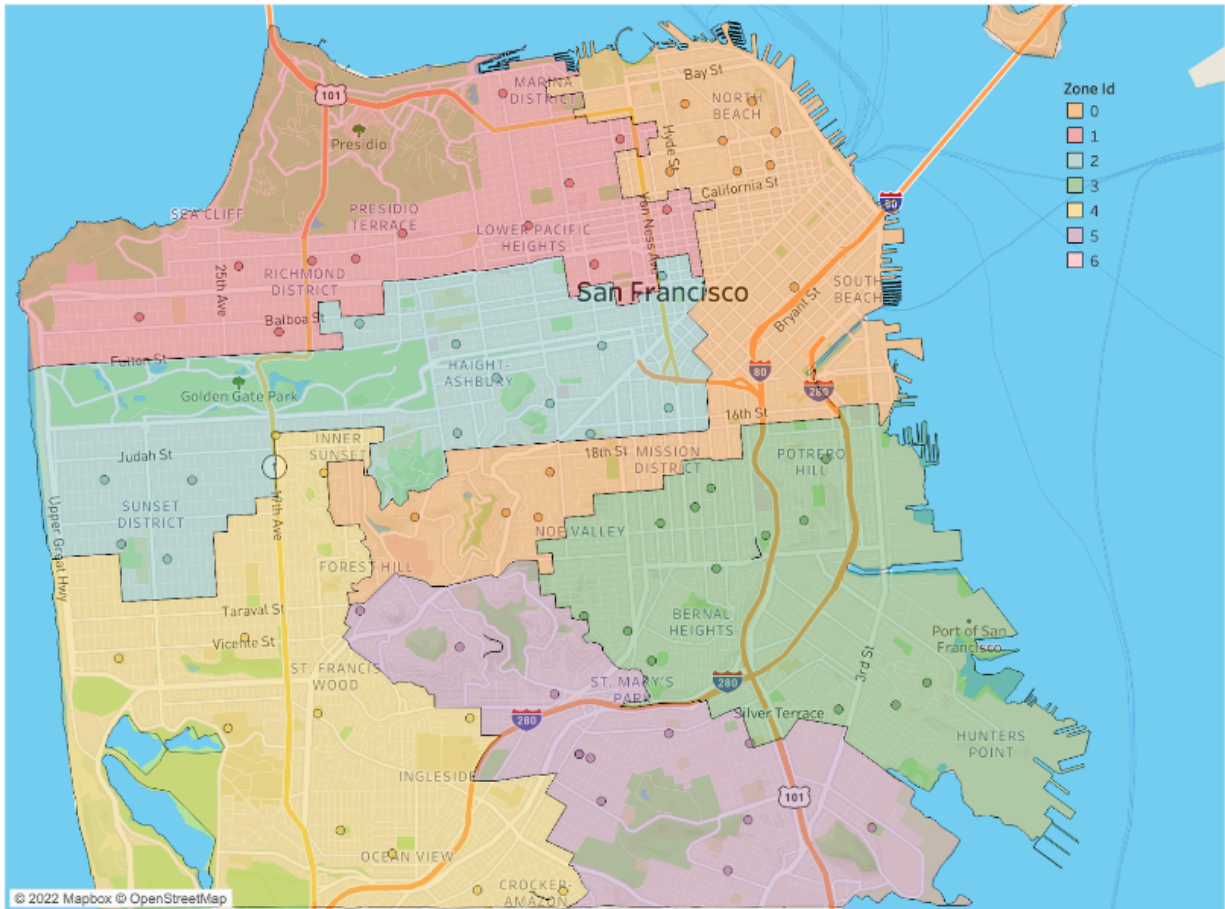
Another concern that arose from the workshop was the lack of aggregation of results. The selection process did not necessarily move stakeholders towards agreement, as each participant could obtain different sets of favorite maps with little overlap. Future iterations of the selection process should be more easily aggregated across users and help participants reach consensus.

Despite these challenges, developing the zone selection tool provided valuable insights into how stakeholders engage with the zone design process and the appropriate level of technical complexity for effective engagement. Future efforts will require significant simplification before expanding to stakeholders less familiar with the problem.

8.4. Focus Groups

1. Look at the shape of the zones.
 - Does anything look good about them?
 - Does anything raise concerns?
 - Are there any zones that really fit the geography of the city (hills, highways, parks, etc.)?
 - Any zones that do not make geographic sense?
2. Using your expertise as a member of your community, neighborhood, or organization, how would the zones affect you and your community?
3. Are there any absolute deal breakers on this map? What do you feel is non-negotiable?
4. What creative solutions can you share for any problems you see?
5. What other feedback do you have about this map? What do you want to see from future maps to better serve students and families?

Leaders provided participants with a sample zone map, and walked through sample discussion prompts to ensure participants understood the task. Then, participants were instructed to look at groupings of 3 zone maps. In addition to the 3 maps, there were brief metrics capturing socioeconomic diversity, travel times, and school performance. Figure 18 shows an example of this map and accompanying metrics. Participants then had unstructured discussion time to investigate the maps. They were encouraged to think out loud, and leaders largely tried to remain out of the discussion, only asking clarifying questions. After roughly 10 minutes of discussion, the group would move on to a second set of maps. The discussion process was repeated 2-3 times, showing participants a total of 6 to 9 maps. The specific maps shown were selected by district staff to represent the broadest possible zone design possibilities. After moving through each group of maps, participants and leaders then regrouped for a final reflection on the zone exploration process.



Map scores out of 10
 The higher the score, the better this map promotes our policy goals

Total map score: 7.1

Diversity score: 7.5
 Predictability score: 7.3
 Proximity score: 5.6

Map Total Score



Table Name

GE_Zoning_2_1_1_2_2_centroids 6-zone-9_BG_Is.shp

Show history

Primary Metrics

Zone Id	% of students who qualify for free or reduced meals	Average home to school walking time (minutes)	Average school performance score (out of 5.00)
0	58.2%	62.33	3.68
1	40.9%	60.97	3.80
2	40.6%	69.42	3.73
3	58.0%	46.44	2.79
4	46.0%	62.80	3.48
5	55.9%	56.16	3.13
6			
Grand Total	50.1%	59.13	3.41



Figure 18: Example map and metrics used in zone design focus groups.

Number of Zones	FRL Balance	Racial Representativeness	Shortage	Initial Count	Final Selection
6	0.10	0.12	0.1	16	No
		0.15	0.1	18	No
	0.15	0.12	0.07	20	Yes
			0.1	14	No
		0.15	0.07	22	Yes
			0.1	20	No
7	0.10	0.15	0.1	8	No
	0.15	0.12	0.7	18	Yes
			0.1	12	No
		0.15	0.7	18	No
	0.1	18	Yes		
10	0.15	0.15	0.15	2	No
		0.2	0.2	6	No
	0.22	0.15	0.1	3	No
			0.15	4	No
		0.20	0.1	11	Yes
			0.15	7	No
13	0.20	0.15	0.3	5	No
		0.2	0.2	4	No
	0.25	0.25	0.12	8	Yes
18	0.22	0.25	0.28	3	No
	0.26	0.25	0.14	5	Yes

Table 2: Constraint selection results for zones that survive the filtering process. The initial count is the number of positions in the consideration set, and the selected count is the positions that were selected via the scoring process.