Automated Redistricting in the American Context: Incorporating Race into Local Search Algorithms

Brian Amos[a]

[a]Department of Political Science, University of Florida, P.O. Box 117325, Gainesville, FL 32611, United States, bamos@ufl.edu

Abstract: The previous literature on automated political districting has addressed many of the generally accepted constraints that an algorithm should contain, such as contiguity, population equality, and compactness. However, the adoption of such methods in the United States has been hindered by the lack of attention paid to a constraint that is legally required in certain contexts: the creation of districts where a majority of the population is of a particular racial minority. This paper will discuss the contexts where such districts are required, and then present three potential methods for including them in a districting objective function. The paper will close with a test case of the least subjective of the three methods, and the results compare favorably to the status quo in place.

1. Introduction

Political districting is a process that seeks to group the subunits of an administrative region into a number of districts, subject to a certain constraints. In some jurisdictions – notably those within the United States – this problem is indeed a political one, where the foremost concerns of those involved are often the fortunes of political parties and current officeholders (Cox & Katz, 2002). For more than 50 years, however, researchers have sought more objective means of districting through the use of models and algorithms. Due to the difficulty of the problem, finding optimal solutions to even simple constraints is a computationally hard problem in real-world scenarios (Altman, 1997), which has led to the exploration of varying heuristic approaches (Ricca, Scozzari, & Simeone, 2011).

One family of heuristic methods is that of the local search algorithm (Ricca & Simeone, 2008). While the specifics of the algorithms differ, at the core of each are a couple of shared features. Firstly, a vector of functions is defined to judge a particular district map on certain objective constraints, and the map that minimizes the resulting score is deemed the best. Secondly, improvements to the map are sought by exploring neighboring solutions: a subunit in one district is migrated into a bordering district, and the objective function is used to judge the change in map quality. The migration may be kept based on the criteria of the method being used, and if so, it becomes the new status quo for the next iteration of the algorithm; otherwise, the old status quo is retained.

As one would expect, then, the quality of the output of one of these algorithms is only as good as the functions being used to score the various constraints. Previous papers have dealt with quantifications of many of the features that have generally been classified as good redistricting practice, such as contiguity, population equality, compactness, and preserving communities of interest (Bozkaya, Erkut, & Laporte, 2003; Bação, Lobo, & Painho, 2005; Yamada, 2009). For the United States, however, a notable gap in the literature has been a function to produce districts, under certain circumstances, where a majority of the population is of a particular racial minority. Due to current federal law and judicial interpretation, this absence is likely to produce illegal maps for many jurisdictions within the country. Given the current practice surrounding redistricting within much of the U.S., this is unfortunate, as it likely would stand to benefit most from the adoption of successful, objective districting methods.

This paper seeks to fill this gap. Section 2 briefly outlines the current state of the law surrounding majority-minority districts in the United States. Section 3 explores some possible methods of incorporating majority-minority districts into an objective function, while Section 4 presents an example of an implementation in action.

2. Race and redistricting in the United States

The 1960s were a decade of major political change in the United States, especially the area of legislative apportionment and districting. Between the end of Reconstruction and up to this point, two trends were present. First was the disenfranchisement of African-American citizens in the South. Despite the ratification of the Fifteenth Amendment to the United States Constitution in 1870, which guaranteed that "The right of citizens of the United States to vote shall not be denied or abridged by the United States or by any State on account of race, color, or previous condition of servitude," the selective use of literacy tests, poll taxes, and outright intimidation by whites in power essentially nullified its effect (Keyssar, 2000). Second was the malapportionment of legislative districts in many of the states, even outside of the South. Congressional districts within a state varied in population by more than 200% in some cases, and state legislative districts could be found that were hundreds of times more populous than the smallest in the state. Quite often, these disparities were a product of making counties or townships indivisible units in districting, which left urban districts – the location of much of the minority population, especially outside of the South – severely underrepresented compared to the rural, primarily white areas (Havard & Beth, 1962). This imbalance was compounded by infrequent redistricting in many states, and the general trend of urban growth outpacing rural.

Residents of these malapportioned districts had no recourse in the federal courts, as the Supreme Court had ruled that districting was a political question, and therefore not justiciable (*Colegrove v. Green*, 1946). Though not strictly a districting case, the first crack in this stand formed in 1960 with *Gomillion v. Lightfoot*, which challenged the change in an Alabama city's boundaries from a square to an odd 28-sided figure, which had the effect of removing nearly all of the African-American voters from the city. This change was struck down on Fifteenth Amendment claims, overturning a District Court ruling that dismissed the case on jurisdictional grounds. 1962's *Baker v. Carr* fully reversed the stand, ruling that political districting cases were indeed justiciable under the powers granted to the Court by Article III of the Constitution and the Fourteenth Amendment's equal protection clause. This opened the door for 1964's *Wesberry v. Sanders* and *Reynolds v. Sims*, where the Supreme Court formally stated its "one person, one vote" doctrine, requiring that district populations within a state needed to be as close in population as practicable. By extension, the cases required the redrawing of districts to balance population every ten years following a U.S. Census.

The requirement for majority-minority districts comes from Section 2 of the Voting Rights Act and the judicial challenges that followed it. In the initial passage of the Act in 1965, Section 2 was essentially a restatement of the Fifteenth Amendment, and was rarely used in subsequent litigation (Boyd & Markman 1983, p. 1352). In 1980, however, the Supreme Court ruled on a Section 2 case in *City of Mobile v. Bolden*. Mobile, Alabama elected the three members of its City Commission in citywide elections rather than by districts, which resulted in African-American candidates being universally unsuccessful in election to the Commission, despite more than a third of the city's population being black. The Court ruled that Section 2 required a discriminatory *intent* for a districting system to be illegal, and since the at-large system had been established in 1911, when African-Americans were effectively barred from holding public office, intent could not be argued, whatever the present effect.

Certain protections in the Voting Rights Act were due to expire in 1982, and in the bill passed that year to extend them was included an amendment to Section 2, written in response to *Mobile v. Bolden*. The original language was designated subsection (a), and the "deny or abridge" language was changed to read "in a manner which results in a denial or abridgement," changing the test from intent to effect. Furthermore, subsection (b) was added, which read in part:

> A violation of subsection (a) of this section is established if, based on the totality of circumstances, it is shown that the political processes leading to nomination or election in the State or political subdivision are not equally open to participation by members of a class of citizens protected by subsection (a) of this section in that its members have less opportunity than other members of the electorate to participate in the political process and to elect representatives of their choice.

This language was tested in 1986 with the case *Thornburg v. Gingles*, where the Supreme Court upheld most of a District Court ruling that struck down several multimember districts in North Carolina. In so doing, the Court established what has become known as "the *Gingles* test" for judging whether illegal racial vote dilution is present. First, "the minority group must be able to demonstrate that it is sufficiently large and geographically compact to constitute a majority in a single-member district." In current practice, this further requires that the majority of the voting-age population (i.e., 18 years and older), not the full population, will be of the minority group. Second, "the minority group must be able to show that it is politically cohesive." Finally, "the minority must be able to demonstrate that the white majority votes sufficiently as a bloc to enable it -- in the absence of special circumstances, such as the minority candidate running unopposed -- usually to defeat the minority's preferred candidate" (p. 51). If all three prongs of the test are met, a majority-minority district is required in order to fulfill Section 2 requirements. In the post-1990 Census redistricting – the first following *Gingles* – the number of House of Representatives majority-minority districts doubled, resulting in a similar jump in the number of African-American and Hispanic members elected (Lublin, 1997).

The first prong of the *Gingles* test, however, is open to interpretation over whether a group is "sufficiently large and geographically compact," and subsequent court cases have led to a complex and occasionally contradictory body of rulings in defining when exactly a majority-minority district is needed. Plans have been struck down for being "geographical monstrosities," with the odd shape of a district seeming only to be explainable in terms of granting majority status to a racial minority (*Miller v. Johnson*, 1996), yet others that have been only marginally better in compactness measures have been upheld. Looking at the reverse, plans have also been struck down on racial vote dilution grounds, prompting more majority-minority districts in the subsequent revisions. While this can lead to confusion and anxiety for a mapmaker looking to follow the law, many of the struck-down cases were extreme examples driven by partisanship. Since African-American and Hispanic voters tend to be Democrats – the former group greatly so – racial issues have been tied up in partisan issues, and mapmakers may attempt to push the boundaries of acceptability to aid their party. An objective function of a local search algorithm, assuming a reasonable balance between race and other factors, seems less likely to run into these problems.

Section 5 of the Voting Rights Act is also worth noting. Unlike Section 2, Section 5 only applies to certain jurisdictions within the country, and those covered must have changes to election laws precleared by either the federal Department of Justice or a federal district court before they are allowed to take effect – this includes changes to district lines. In a jurisdiction's defense of a new districting plan, it must show that the new law does not have the intent to dilute minority votes, and does not have the effect of diminishing the ability of a minority to elect candidates of their choice compared to the previous plan. In practice, this often means that a new plan cannot have fewer majority-minority districts than the one it is replacing, unless there is justification on the grounds of demographic shifts over the preceding decade. At the time of this writing, Section 5 has lost its relevance due to the 2013 Supreme Court case *Shelby County v. Holder*, which struck down the portion of the Voting Rights Act that defined the formula for which jurisdictions are covered under Section 5. As such, the section no longer applies anywhere; Congress has the ability to pass a new formula, but given the current political climate, this may not happen soon.

3. Incorporating majority-minority districts in local search algorithms

There are several means by which majority-minority districts can be produced in automated legislative redistricting. This paper will discuss three options, which primarily vary on how much of the decision making of where majority-minority districts should be drawn is done in advance of running the algorithm. At the root of the issue is whether all African-American voters (or Hispanic, or any other minority) within a jurisdiction are similar enough for a mapmaker to be unconcerned with which ones in particular get grouped together to form a majority within a district. The Supreme Court in *Shaw v. Reno* warned that, "A reapportionment plan that includes in one district individuals who belong to the same race, but who are otherwise widely separated by geographical and political boundaries, and who may have little in common with one another but the color of their skin, bears an uncomfortable resemblance to political apartheid" (1993, p. 647). Most algorithms found in the academic literature do weight for compactness and following existing political boundaries, but even if a given majority-minority district is legal, there can still be normative concerns: the district can represent a grouping of residents that do not see themselves as a single group. Miami, Florida has a large Cuban population, but also a mix of Hispanic voters with other national heritages; for the purposes of the Voting Rights Act, they are considered a unified group, but there can be large cultural and political differences among the subgroups (Scher, Mills, & Hotaling, 1997, p. 160). For relatively large districts, such as those for the House of Representatives, it may not be possible to make such distinctions and still create majority-minority districts, but for districts on the scale of a city government, it may be a concern worth addressing.

At the same time, though, one of the major strengths of automated redistricting is the reduction of subjectivity in placing lines. Some cities exhibit high levels of racial segregation, but in most cases, the edges of a racial community will be blurry, requiring a judgment call if a mapmaker chooses to define them in advance. Given the partisan leanings discussed above that minority groups tend to have, this sort of leeway may result in legislative fights as bitter as if automated methods were not used at all.

Three potential methods for including majority-minority districts in an objective function will be explored below. As stated in the introduction, the objective function gives a particular district map a score based on several terms – one for each constraint being observed – that are summed together; the map that produces the lowest score is considered the best. The three majority-minority district terms that will be presented, going from the most to least subjective, will be referred to as *geographic area*, *target range*, and *borderline acceleration*.

3.1. Geographic area

The redistricting practice of preserving communities of interest can have strong similarities to drawing majority-minority districts, and the former has been addressed in the work done by Burcin Bozkaya and co-authors (Bozkaya 1999; Bozkaya et al., 2003). The term “community of interest” is notoriously slippery, with observers complaining that it can “mean almost anything one chooses” (Arrington, 2010, p. 6); even between U.S. states, the legal definition can vary. In general, though, a community of interest is a group of people with a shared social, cultural or economic identity that raises unique legislative needs. One can imagine relatively diffuse groups that meet this definition. For example, the gay and lesbian community would seem to fit well, but with a few notable exceptions, such as the Castro District in San

Francisco or the West Village in New York City, they are not segregated into distinct neighborhoods.

However, the previous work in automated redistricting on communities of interest has been on those that can be defined as geographic areas, and this may extend well to racial majority-minority districts in certain cases. Racial communities often are considered a type of community of interest themselves, and often are relatively segregated so as to make a geographic definition possible. Bozkaya's work has been in city districts, and worked under the assumption that each subunit in the entire jurisdiction could be assigned to a community or neighborhood. As such, communities are treated in roughly the same manner as when changes from the previous district map are sought to be minimized, with an eye towards minimizing splits. The function is given as

$$f_{mm}(x) = 1 - \sum_{j \in J} G_j(x) \Big/ \sum_{j \in J} P_j(x)$$

where J is the set of districts in solution x, $G_j(x)$ is the population of the largest community in district j, and $P_j(x)$ is the total population of district j. Since racial communities are unlikely to extend across the entire jurisdiction, the remaining area not covered can be assigned to a single "leftover" community, and this formulation will continue to be effective.

As stated before, this method can be hindered when the boundaries of a racial community are not clearly defined; judgment calls would need to be made, removing one of the benefits of objective, automated districting. Even when community lines are clear, though, there is an additional wrinkle that is not present for normal communities of interest: the community must be large enough to constitute the majority of the voting-age population in a district. If a particular racial group is separated into several communities for the purposes of the algorithm, and each is under the population level needed to reach a majority, it is possible to have an end result with no majority-minority districts. If two or more of these communities were near enough to have been combined to reach a majority without sacrificing too much on compactness, the resulting map may be deemed a violation of Section 2 of the Voting Rights Act.

Thus, for the purposes of producing majority-minority districts, there is no reason to define a racial community that is below the threshold of being the majority of a district – one may wish to include a communities-of-interest term within the objective function and include a smaller racial community there, however. In especially large, dense racial communities, this may not be a problem, but when the minority population is more diffuse, the geographic approach can be nearly equivalent to drawing the majority-minority district in advance.

3.2. Target Range

The previous approach used predefined racial communities to generate majority-minority districts, requiring the mapmakers to look at racial data and convert them into a simplified picture of the jurisdiction. The following two approaches remove this before-the-fact process and look at the racial statistics of each subunit directly. As such, these methods require that the source database includes the full voting-age population for each subunit, as well as the breakdown for each racial minority that could potentially form a majority-minority district.

Altman and McDonald (2011) do make note of majority-minority districts in the paper describing their open-source automated redistricting software, which includes local search algorithms for plan refinement. Though they do not formally express how majority-minority districts can be included in an objective function in their paper, one can look at the source code (included as online supplemental material with their paper) to find out how they score for race. A target range is supplied for the proportion of a district that is of a particular minority – generally, this will be a range with a minimum endpoint larger than 0.5, in order to guarantee majority-minority status. If a district falls within the range, it is given a score of zero, and the score increases linearly as the proportion moves above or below the target. To be most effective, a certain subset of districts should be chosen to be scored based on the majority-minority district term, while the rest are not. Otherwise, the movement of a dense-minority subunit from a district with a low overall racial proportion to a district approaching the target range – a trade that is desirable – will have roughly zero net effect on the objective function score, since the decrease in proportion of the former district would cancel out the increase in proportion of the latter.

A formal definition of the term, slightly adapted from Altman and McDonald, is given as

$$f_{mm}(x) = \frac{\sum_{k \in K} T_k(x)}{n}, \qquad where$$

$$T_k(x) = \begin{cases} 0, & if\ min \le R_k(x) \le max \\ R_k(x) - max, & if\ R_k(x) > max \\ min - R_k(x), & if\ R_k(x) < min \end{cases}$$

where *K* is the subset of *n* districts assigned to be majority-minority districts in solution *x*, *max* and *min* are the endpoints of the target range, and $R_k(x)$ is the proportion of voting-age residents in district *k* of the particular minority the district was assigned. A plot of one district's contribution to this term, using the default range supplied by Altman and McDonald (0.65 to 0.7), is given in Figure 1.

[FIGURE 1 ABOUT HERE]

This method is effective in creating majority-minority districts, and has the benefit of not being "greedy": over a certain racial proportion, a district is punished for adding more minority residents. Though this method removes some subjectivity by not requiring mapmakers to identify minority communities in advance, it still requires a mapmaker to choose how many majority-minority districts to aim for, and to choose which districts in the seed map will be the ones that strive for majority-minority status. As discussed above, majority-minority districts are usually equivalent to packed Democratic districts, making the choice of the number of these districts a politically charged decision.

3.3. Borderline accerleration

Instead of an *ex ante* identification of which districts are to become majority-minority districts, a more objective approach would be to design a function that had three properties: first, it would encourage districts near the threshold of majority-minority status to add and shed subunits to cross that threshold, second, retain majority-minority districts through the iterations of the algorithm, and third, be indifferent towards the racial composition of districts well above and below the majority-minority threshold. The inverse logit function, when properly scaled, meets these requirements. A formulation of the function for a local search algorithm, which is reversed to make minimization the goal, can be given as

$$f_{mm}(x) = \frac{\sum_{j \in J} \frac{1}{e^{(\alpha Rj(x) - \beta)} + 1}}{n}$$

where *J* is the set of *n* districts in solution *x*, $R_j(x)$ is the proportion of the voting-age population belonging to a certain minority in district *j*, and α and β are pre-defined constants that adjust the function's behavior. Since the function only applies to one race, there will need to be a separate term for each race that has the potential of forming a majority-minority district. A plot of one district's contribution to this term, using α = 25 and β = 11, is given in Figure 1.

[FIGURE 2 ABOUT HERE]

Below a proportion of 0.2 and above 0.7, minor changes in the racial composition of a district have essentially no impact on the objective score. This is desired: those districts in the former group are unlikely to become majority-minority districts, so their racial composition should have little weight in deciding whether a change is beneficial, and those in the latter group are safely majority-minority districts, so it would be harmful to neighboring potential majority-minority districts for them to continue to be greedy. It is only when a district falls between roughly 0.35 and 0.55 that changes will have a large effect on the decision to maintain or reject a change, and it encourages potential majority-minority districts to reach a safe, but not overwhelming racial proportion. The α and β constants can be adjusted to tweak the proportions where changes become important; increasing α steepens the drop in the graph, while increasing β moves the graph to the right.

A drawback to the method is that it can produce an unpredictable number of majority-minority districts. For some jurisdictions, such as those covered under Section 5 of the Voting Rights Act, being able to set a target number of districts may be desirable. However, the major strength is that it keeps best with the spirit of automated redistricting. It requires no subjective input from a mapmaker in advance, and everything is driven by Census data. Given the long history of gerrymandering in the United States, providing room for manipulation may undermine the trust in a system that is otherwise objective, and make the argument for adoption weaker. To show the effectiveness of the inverse logit method, the following section will provide a test case of it in action.

4. Test case: Arizona State Legislature

Thanks to a citizen-initiated ballot measure, Arizona is one of the few states that have put the task of state and congressional redistricting into the hands of an independent commission, rather than politicians. Without partisanship being the driving force behind the maps, the status quo is a stricter standard to compare against than one in a gerrymandered state. This test case will draw a thirty-district map for the Arizona State Legislature – both houses use the same districts, but elect a different number of representatives from each – using the data available from the 2010 U.S. Census.

Arizona is located in the southwestern United States, and borders Mexico. As such, it has a relatively large Hispanic population, making up 25% of the voting-age population as of the 2010 Census. The remaining non-white population is relatively small: African-Americans and American Indians are the next two largest minority groups, each comprising about 4% of the state. The African-American population is diffuse, though; Phoenix is the only city with a population greater than 100,000 where the African-American proportion exceeds 5%, and furthermore, no zip code within the city has a proportion greater than 25%. This makes an African-American majority-minority district unlikely, and in testing the algorithm in various ways, no district ever came close to reaching majority status. On the other hand, the American Indian population is mostly found on the large Navajo and Hopi reservations in the northeastern part of the state, making it possible for a majority-minority district centered there to form.

The subunit used for this test case was the Census tract, a sub-division of the county used by the U.S. Census; this geography was chosen due to the constitutional requirement to not split tracts "to the extent practicable" (Arizona Constitution). Arizona had 1526 Census tracts in 2010, resulting in an average of about 51 per district. All of the geographic and demographic data used comes from the U.S. Census Bureau and were obtained from the Census's FTP server. The demographic data, specifically, come from the Census's 2010 Redistricting File, a dataset released in early 2011 in advance of the full Summary File to give states more time to handle the redistricting process, and contains only population and racial data.

4.1. The Old Bachelor Acceptance algorithm

Though the proposed method of generating majority-minority districts can be used in other types of local search algorithms, I've chosen to follow the suggestion of Ricca and Simeone and their comparative tests (2008) and use Old Bachelor Acceptance, an algorithm first suggested by Hu, Kahng, and Tsao (1995). Both papers can be referenced for a detailed description of the method, but a brief overview is worthwhile. The process begins with a seed district map that is provided by the mapmaker. Each iteration of the algorithm selects a random district within the map, and then temporarily changes a neighboring subunit to be a member of the district. The objective function is then used to score this new map, and the change in quality from the old is noted. If the change is an improvement, or causes a worsening within a certain threshold, it is made permanent; otherwise, the subunit is returned to its original district. The novel part of the algorithm is that in the case of a change being kept, the threshold for retaining a drop in map quality is decreased, and when a change is rejected, the threshold is increased. This creates a repeating cycle of searching and optimization, allowing for a broader range of district configurations to arise and be tested. The map with the lowest objective score over the course of the algorithm is deemed the final map.

4.2. The objective function

Two factors are hard-coded into the algorithm and are not part of the objective function. First, any change that would cause a district to become non-contiguous is rejected automatically. Second, though population equality is a term within the function, a legal rule-of-thumb for state legislative maps is that the deviation in district populations cannot exceed 10% of the average district population; the true case law is more complex, allowing deviations above this figure in some cases and rejecting those below in others, but 10% is a generally accepted guideline (Levitt, 2010, p. 44). Thus, any change that would make a district's population 5% above or below the target population of a district is also rejected automatically. The rest of the factors considered are included in the objective function, which is given as

$$F(x) = \gamma_{comp} f_{comp}(x) + \gamma_{popeq} f_{popeq}(x) + \gamma_{county} f_{county}(x) + \gamma_{mm} f_{mm}(x)$$

where the γ coefficients are weights for each term, $f_{mm}(x)$ is as defined in Section 3.3, while the remaining terms are as defined below.

4.2.1. District compactness

There are many ways to measure district compactness, and no measure by itself is perfect in capturing all the meanings people associate with the term "compact" (Niemi, Grofman, Carlucci, & Hofeller, 1990). For this test case, the measure used will be a perimeter-area measure, which highly values districts that minimize their perimeter lengths relative to their areas – a circle would be the perfect shape in this case. This is given with the formula

$$f_{comp}(x) = \frac{\sum_{j \in J} 1 - \frac{4\pi A_j(x)}{\left(M_j(x)\right)^2}}{n}$$

where J is the set of n districts in solution x, and $A_j(x)$ and $M_j(x)$ are the area and perimeter of district j, respectively.

4.2.2. Population equality

As said above, there is a hard-coded upper and lower limit to the amount of population deviation allowed, but all else being equal, a closer adherence to the principle of "one person, one vote" is better. Since there can be no deviation greater than 5%, this figure is used for normalization.

$$f_{popeq}(x) = \frac{\sum_{j \in J} \left| P_j(x) - t \right|}{0.05t * n}$$

where J is the set of n districts in solution x, $P_j(x)$ is the population of district j, and t is the target population for a district, defined as $\frac{\sum_{j \in J} P_j(x)}{n}$ rounded to the nearest integer.

4.2.3. Respecting county boundaries

A common redistricting goal is to minimize the number of counties split across districts, and vice versa. The index used comes from Ricca and Simeone, and is defined as "the percentage of those [subunits] that have at least one adjacent unit belonging to the same electoral district, but to a different administrative area" (2008, p. 1412). Formally,

$$f_{county}(x) = \frac{\sum_{j \in J} \frac{S_j(x)}{C_j(x)}}{n}$$

where *J* is the set of *n* districts in solution *x*, $S_j(x)$ is the count of subunits in district *j* that have an adjacent subunit in *j* in a different county, and $C_j(x)$ is the total count of subunits in district *j*. Since the test case is using Census tracts as a subunit, which can cut across municipal boundaries, it is not possible to have a term that respects city lines. However, with a smaller unit of geography this could be achieved, and the formulation would be similar to the county term.

4.3. Seed map

Since local search algorithms optimize an existing map, one must be provided. While the previous decade's state legislative map was an option, a random approach was adopted instead, to keep the entire process automated from start to finish. The method chosen to draw the seed plan was William Vickrey's multi-kernel growth procedure (Vickrey, 1961), as it creates relatively compact and even-population districts. Briefly, a random subunit *A* is chosen. The farthest unassigned subunit from *A*, subunit *B*, is found. Then, the closest unassigned subunits to *B* are grouped together as a district until the target district population is met, at which point the process repeats, until all districts are assigned. The algorithm as just described is not perfect: it is not uncommon for an unassigned pocket of tracts to become surrounded by assigned tracts, so that it is unable to grow to the population of a full district. Thus, corrections are sometimes necessary, where the unassigned pocket is given to an adjacent existing district, and a portion of that district adjacent to a larger unassigned region is reset to unassigned.

4.4. Data preparation

The relevant Census demographic information was collected for each subunit and added to a database for use in the algorithm. Then, the neighboring subunits for each subunit were found; units that only met at single points, rather than sharing an edge, were not considered adjacent, and such relationships were not included in the neighbor table. Finally, geographic preprocessing took place. Areas for each subunit were calculated in ArcGIS, as were the vertices and perimeters for the seed map districts.

4.5. Results

Twenty seed maps were created, and the Old Bachelor Acceptance algorithm as described was used on them all. The map that minimized the objective function is presented in Figure 3, alongside the actual map in effect for Arizona. Both include county lines for reference. Additionally, Table 1 includes some statistics comparing the two maps.

[FIGURE 3 ABOUT HERE]

[TABLE 1 ABOUT HERE]

Each of the twenty maps produced one American Indian majority-minority district. Due to the location of the reservations in the corner of the state and the way by which Vickrey's method produces districts, many of the seed maps had this majority-minority district from the beginning. The number of Hispanic districts ranged from five to seven in the final maps, up from between two and five in the seed maps. As Table 1 shows, both the best test case map and the independent commission map had one American Indian majority-minority district and seven Hispanic, counts which roughly reach levels proportional to the populations within the state. However, the test case map is an improvement on the factors of population equality, the number of counties split by two or more districts, and compactness as defined for the objective function (i.e., lower values are more compact). Granted, the Arizona independent commission took into account more constraints than the ones in Table 1, as it was required to preserve communities of interest and encourage competitive districts where it did not conflict with the other requirements (Arizona Constitution). Even considering that fact, though, finding improvements across all other factors was unexpected, and helps prove the strength of local search algorithms and the success of the majority-minority district term.

5. Conclusion

Over the past fifty years, the academic literature on automated districting has seen great leaps in the quality of methods presented and the resulting maps produced. Coupled with the leaps made in computing technology, an affordable home computer is now able to produce districts for a jurisdiction subject to most of the major constraints lauded by good-government groups. The goal of this paper was to fill a weakness in the potential success of importing the methods to the American context, that of the ability to create racial majority-minority districts. As shown in the test case, the method presented was successful in producing majority-minority districts in an objective manner while still balancing other, traditional redistricting constraints.

There are certainly some troubles in adopting automated redistricting, especially a heuristic method like the one used in this paper (Altman 1997; Altman & McDonald, 2010). Choosing which terms to include in an objective function and the relative weights between them always involves some subjectivity, and even when an agreement can be reached on that front, heuristic methods can never guarantee that the result produced is the best possible. With that said, plans created by automated methods have been adopted in places (e.g., Edmonton, Alberta, Canada's city wards; see Bozkaya et al., 2011), and may have a chance in the United States within smaller-scale jurisdictions, like county or city commissions. Even if the methods are never adopted, though, the maps produced can provide a valuable benchmark with which to compare the actual plans under consideration during the redistricting process, or can even be used as a starting point for mapmakers. By allowing for the consideration of race, this paper serves to make these possibilities more realistic than before.

References

Altman, M. (1997). The Computational Complexity of Automated Redistricting: Is Automation the Answer? *Rutgers Computer and Technology Law Journal*, *23*(1), 81-142.

Altman, M. & McDonald, M. P. (2010). The Promise and Perils of Computers in Redistricting. *Duke Journal of Constitutional Law and Public Policy*, *5*(1), 69-112.

Altman, M. & McDonald, M. P. (2011). BARD: Better Automated Redistricting. *Journal of Statistical Software*, *42*(5), 1-28. http://www.jstatsoft.org/v42/i04.

Arizona Constitution. Article IV, part 2, § 1(14)-(16).

Arrington, T. S. (2010). Redistricting in the U.S.: A Review of Scholarship and Plan for Future Research. *The Forum, 8*(2). doi:10.2202/1540-8884.1351

Bação, F., Lobo, V., & Painho, M. (2005). Applying genetic algorithms to zone design. *Soft Computing*, *9*(5), 341-348.

Baker v. Carr, 369 U.S. 186 (1962).

Boyd, T. M. & Markman, S. J. (1983). The 1982 Amendments to the Voting Rights Act: A Legislative History. *Washington and Lee Law Review*, *40*(4), 1347-1428.

Bozkaya, B. (1999). *Political districting: A tabu search algorithm and geographical interfaces.* (Doctoral dissertation). Retrieved from ProQuest Dissertations and Theses. (Order No. NQ39508, University of Alberta (Canada)).

Bozkaya, B., Erkut, E., & Laporte, G. (2003). A tabu search heuristic and adaptive memory procedure for political districting. *European Journal of Operational Research*, *144*(1), 12-26.

Bozkaya, B., Erkut, E., Haight, D., & Gilbert Laporte. (2011). Designing New Electoral Districts for the City of Edmonton. *Interfaces 41*(6), 534-547.

City of Mobile v. Bolden, 446 U.S. 55 (1980).

Colegrove v. Green, 328 U.S. 549 (1946).

Cox, G. W. & Katz, J. N. (2002). *Elbridge Gerry's salamander: The electoral consequences of the reapportionment revolution*. Cambridge: Cambridge University.

Gomillion v. Lightfoot, 364 U.S. 339 (1960).

Havard, W. C. and Beth, L. P. (1962). *The politics of mis-representation*. Baton Rouge: Louisiana State.

Hu, T. C., Kahng, A. B. and Tsao, C. A. (1995). Old Bachelor Acceptance: A New Class of Non-Monotone Threshold Accepting Methods. *ORSA Journal on Computing*, *7*(4), 417-426.

Levitt, J. (2010). *A citizen's guide to redistricting*. New York: Brennan Center.

Lublin, D. (1997). *The paradox of representation: Racial gerrymandering and minority interests in Congress.* Princeton.

Keyssar, A. (2000). *The right to vote: The contested history of democracy in the United States*. New York: Basic Books.

Miller v. Johnson, 515 U.S. 900 (1995).

Niemi, R. G., Grofman, B., Carlucci, C., & Hofeller, T. (1990). Measuring Compactness and the Role of a Compactness Standard in a Test for Partisan and Racial Gerrymandering. *Journal of Politics*, *52*(4), 1155-1181.

Reynolds v. Sims, 377 U.S. 533 (1964).

Ricca, F., Scozzari, A., & Simeone, B. (2011). Political districting: from classical models to recent approaches. *4OR, 9*(3), 223-254.

Ricca, F. & Simeone, B. (2008). Local search algorithms for political districting. *European Journal of Operational Research, 189*(3), 1409-1426.

Scher, R. K., Mills, J. L., & Hotaling, J. J. (1997). *Voting rights and democracy: The law and politics of districting*. Chicago: Nelson Hall.

Shaw v. Reno, 509 U.S. 630 (1993).

Shelby County v. Holder, 570 U.S. (2013).

Thornburg v. Gingles, 478 U.S. 30 (1986).

Vickrey, W. (1961). On the Prevention of Gerrymandering. *Political Science Quarterly*, *76*(1), 105-110.

Voting Rights Act, 42 U.S.C. §1973-1973aa-6 (1976).

Voting Rights Act Amendments of 1982, Pub. L. No. 97-205, 96 Stat. 131 (1982).

Wesberry v. Sanders, 376 U.S. 1 (1964).

Yamada, Takeo. 2009. A mini-max spanning forest approach to the political districting problem. *International Journal of Systems Science*, *40*(5), 471-477.

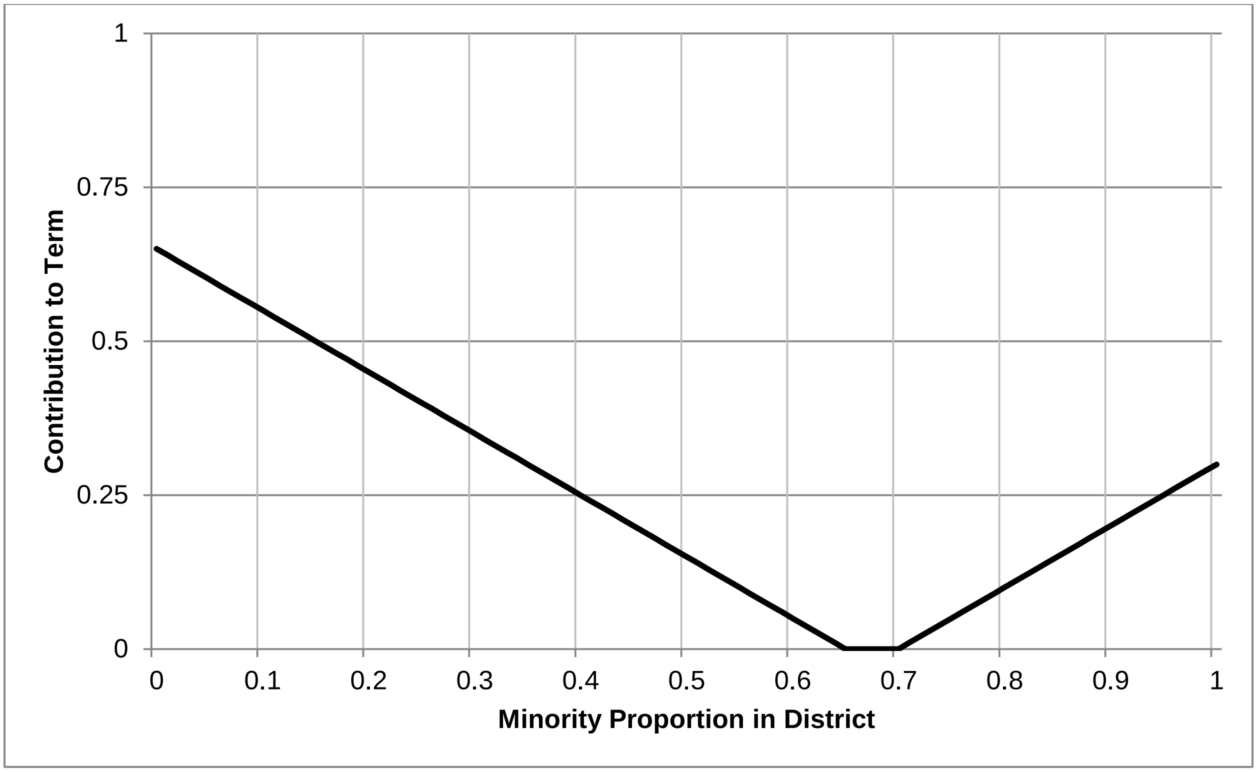

Figure 1: Single district contribution to the target range majority-minority objective function term when the range is (0.65, 0.7).

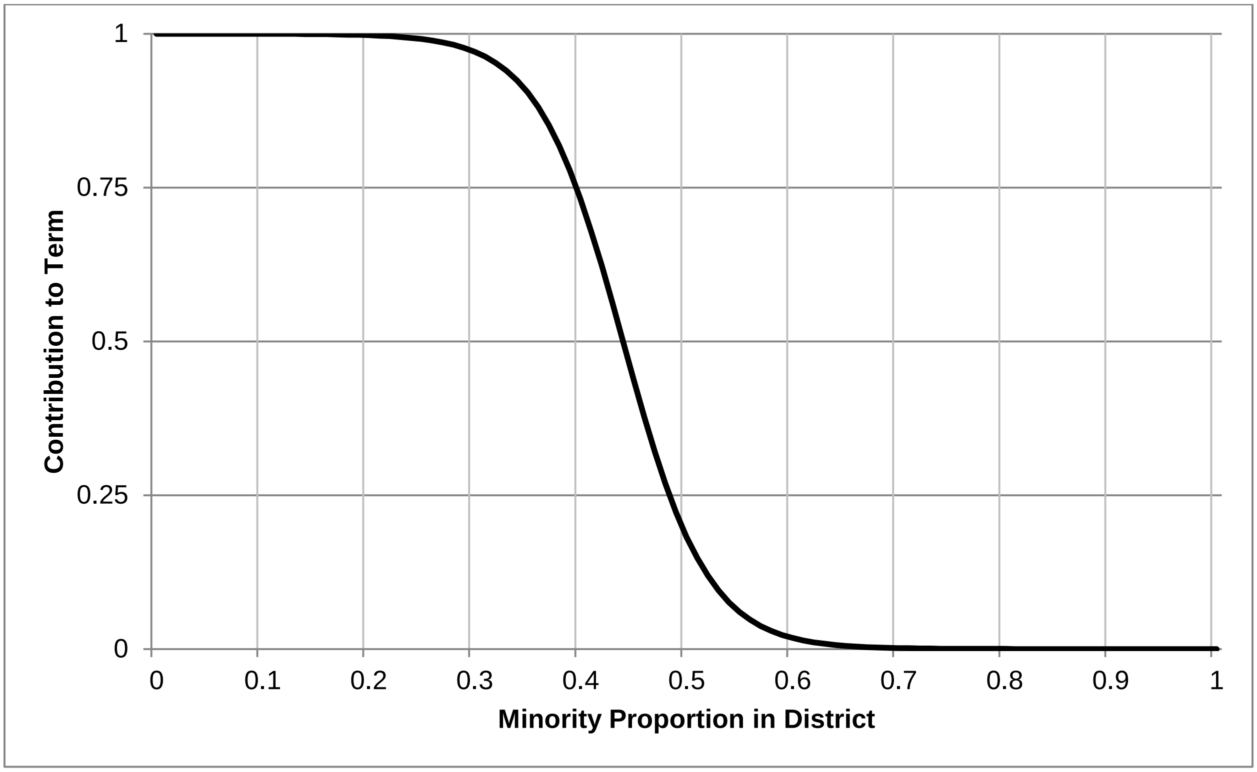

Figure 2: Single district contribution to the majority-minority objective function term when α = 25 and β = 11.

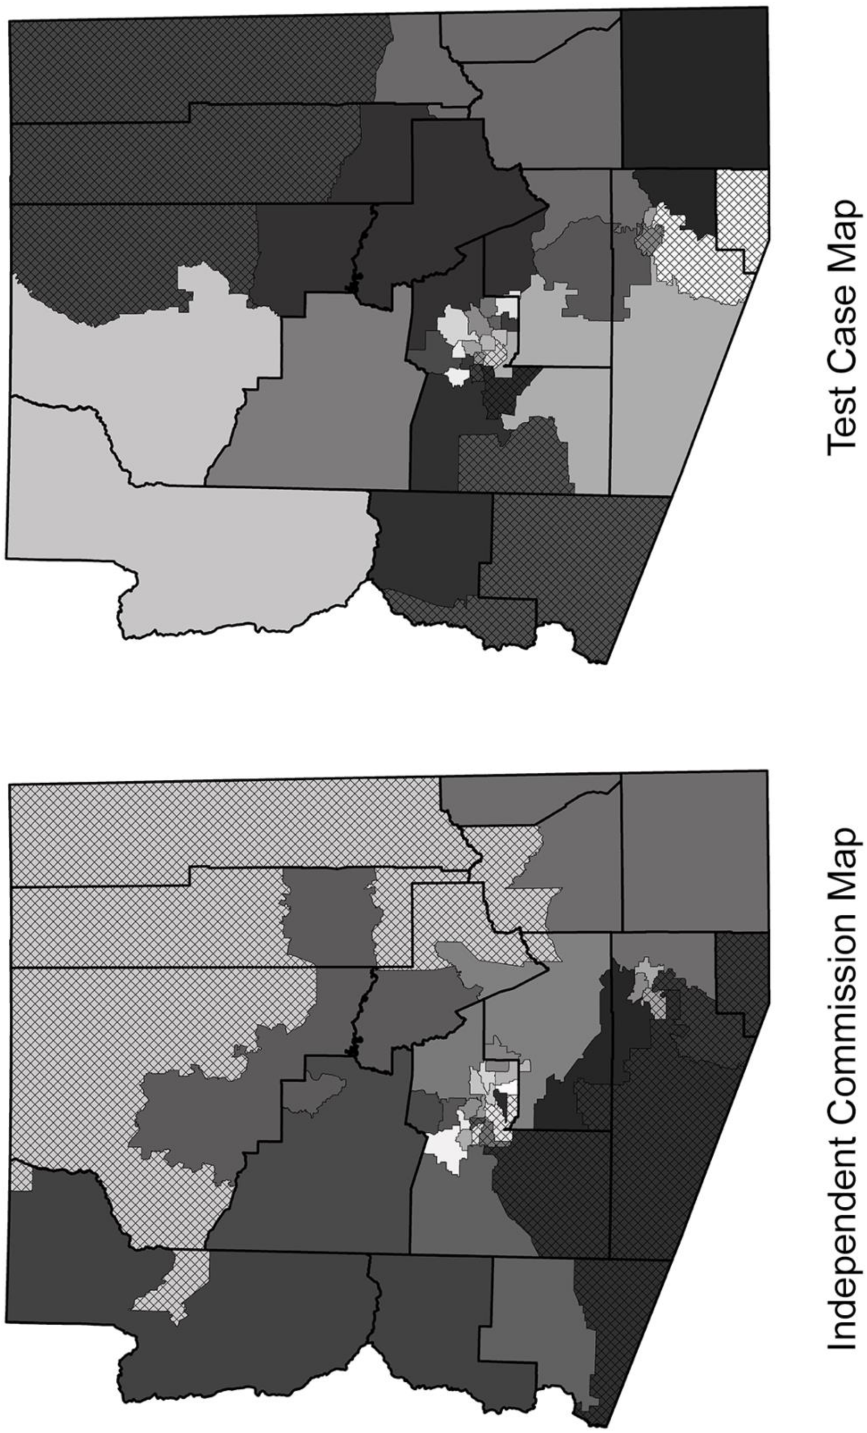

Figure 3: Comparison of Arizona State Legislature maps. Crosshatched districts are majority-minority

	Independent Commission Map	**Test Case Map**
Number of American Indian majority-minority districts	1	1
Number of Hispanic majority-minority districts	7	7
Population deviation	8.8%	7.7%
Split counties (out of 15)	10	7
Average district area-perimeter compactness	0.694	0.551

Table 1: Comparison of Arizona State Legislature maps