Published in final edited form as:

Urban Geogr. 2011 ; 32(3): 334–359. doi:10.2747/0272-3638.32.3.334.

Identifying and Bounding Ethnic Neighborhoods

John R. Logan,
Brown University

Seth Spielman,
University of Colorado – Boulder

Hongwei Xu, and
Brown University

Philip N. Klein
Brown University

Abstract

This study presents three novel approaches to the question of how best to identify ethnic neighborhoods (or more generally, neighborhoods defined any aspect of their population composition) and to define their boundaries. It takes advantage of unusual data on the residential locations of all residents of Newark, NJ, in 1880 to avoid having to accept arbitrary administrative units (like census tracts) as the building blocks of neighborhoods. For theoretical reasons the street segment is chosen as the basic unit of analysis. All three methods use information on the ethnic composition of buildings or street segments and the ethnicity of their neighbors. One approach is a variation of k-functions calculated for each adult resident, which are then subjected to a cluster analysis to detect discrete patterns. The second is an application of an energy minimization algorithm commonly used to enhance digital images. The third is a Bayesian approach previously used to study county-level disability data. Results of all three methods depend on decisions about technical procedures and criteria that are made by the investigator. Resulting maps are roughly similar, but there is no one best solution. We conclude that researchers should continue to seek alternative methods, and that the preferred method depends on how one's conceptualization of neighborhoods matches the empirical approach.

North American cities have long been spatially segmented by the race, ethnicity, social class, and nativity of their residents. Consequently when social scientists write about a neighborhood, we can scarcely resist adding an adjective that connotes its population composition. We think of middle class or working class or gentrifying neighborhoods, immigrant neighborhoods or more specifically Latino or Chinese neighborhoods, and white or black neighborhoods as though the neighborhood were defined by these characteristics. Even when we have in mind socially mixed neighborhoods or neighborhoods in transition, these categories gain meaning from the implied reference to neighborhoods that are not mixed, or not changing, or the way the neighborhood was or is going to become.

We review three approaches to identifying neighborhoods and their boundaries based on information about their social composition. As we show, each of these methods is rooted in a distinct set of assumptions about the nature of neighborhoods. We exploit an unusual data file that includes the geocoded address of every household in Newark, NJ, in 1880, along with personal characteristics of residents as recorded in the 1880 Census of Population. Such information is available for other populations, such as the contemporary registration data for

residents of several North European countries and sometimes census data (such as the Israeli census data analyzed by Omer and Or [2005] or U. S. data at a confidential Census Research Data Center). In this case the data are fully public, and there are no limitations on how it can be analyzed. We address the question of how methods appropriate for individual-level address data can be applied to the study of ethnic neighborhoods.

Three very different approaches are described and applied to the Newark case. The first is a variation of k-functions, calculating the density of same-group and other-group population at various distances from each resident. The second is a data mining technique commonly applied to the problem of sharpening photographic images, where decisions are made about whether a pixel of a given color is “really” that color or should be transformed to the color of its neighbors. The third is a Bayesian model that has been used to detect spatial clusters. While these approaches involve widely diverging procedures, they have two major elements in common. They all take account of information about people and their neighbors at varying distances, presuming that an “ethnic neighborhood” is one where there is a disproportionate presence of members of a particular group within some local area. In addition, they all result in spatially coherent zones of the city that are classified as ethnic (and assigned to a specific ethnic group) or “mixed.”

Our problem of identifying neighborhoods is related to geographers’ longstanding interest in the definition of areal units for spatial analysis. Neighborhoods are geographic regions, but they are also collections of people, institutions, and infrastructure. Because we begin with data about individual people our problem is very much like the one that Hartshorne (1939, Chapter 11) enunciated decades ago: “We could study the geography of the area only from the study of the geography of the infinite number of points within it. This task, being infinite, is impossible. The problem of regional geography, as distinct from a geography of points, is how to study and present the geography of finite areas.” The task of combining points (or other units) into coherent finite areas is called regionalization. Early proponents of regionalization like Openshaw (1977) saw the approach as a way to solve the Modifiable Areal Unit Problem, which is that the results of statistical analyses change as geographic units of analysis change in size and shape (Fotheringham and Wong 1991). King (1997) argues that the MAUP arises only when we don’t know the appropriate unit of analysis on theoretical grounds. Our view is that the issues are both theoretical and empirical. Our problem is not how to select from among a number of pre-formed alternatives (such as census blocks, tracts, precincts, cities, or counties) based on knowing which is theoretically “right.” We must become engaged in creating the appropriate units, using empirical tools like those applied to regionalization in a way that also incorporates varying conceptualizations of neighborhoods.

Approaches to identifying ethnic neighborhoods

Recent ethnographic studies of cognitive maps show that residents draw on many of the same characteristics to identify the boundaries of their neighborhoods as do social scientists. Lacy (2007) describes this as “boundary work,” and focuses on how residents of middle class black areas seek to signify that their neighborhood is distinct from neighborhoods of black working class or poor people. In a Baltimore neighborhood Rich (2009) found that white residents used both the racial composition and class background of specific blocks to mark the limits of their own neighborhood: “when speaking about the boundaries for themselves, many talked about race and, in the terms of one Villager, where the racial ‘Magnet Line’ fell. Homeowners deemed the more middle-class and white areas of the neighborhood ‘the heart of Village Heights,’ while they questioned if the areas with a majority of black residents were included in the boundaries” (p. 838). But in another locale Campbell et al (2009) reported that white residents preferred to draw wider boundaries in

order to think of their neighborhood as more racially and occupationally diverse: they drew “large maps that extended beyond the census tract to include a northern area. Whereas the Broadmore block group was a white, middle to upper middle class neighborhood, North Broadmore was less affluent and included a higher percentage of African Americans” (p. 479).

Neighborhoods often have names. They sometimes have clear boundaries, and Suttles (1972, p. 4, cited in Stoneall 1981) argues that residents tend to construct simplified images of the city in which differences between neighborhoods are magnified. Simplified images, he believes, “serve us well by reducing the complexity of the urban landscape to a range of discrete and contrastively defined ecological units despite the general continuity, gray areas, and constant changes in any section of the city. A cognitive map of our urban environs is useful for precisely the reason that it simplifies to the point of exaggerating the sharpness of boundaries, population, composition, and neighborhood identity.” Nonetheless researchers have noticed that it is common for borders to be fluid. Hunter (1974), for example, reported that areas of Chicago that he studied had “rolling” boundaries – people might agree on the name of their neighborhood, but those living near its edge tended to perceive it as extending further in that direction. Ambiguous boundaries probably work fine for residents, but they complicate neighborhood research. How can we measure neighborhood characteristics and their changes over time if we only know vaguely where the neighborhoods are?

If neighborhoods exist in the minds of residents, then the direct approach is to find out how residents understand their neighborhoods. In fact, as Logan, Alba and Zhang (2002, pp. 303–304) point out, “[E]thnic neighborhoods are most often identified and studied through fieldwork, where the researcher typically begins with the knowledge that the ethnic character of a given locale is socially recognized – certainly by group members and perhaps also by others. This ethnic character may be visible through observation of people in public places, the names of shops or the languages found on signs or spoken by clerks or patrons, or by community institutions such as churches, social clubs, and associations.” Yet many researchers are unable to do original fieldwork, and many study designs require information on more locations than could feasibly be covered by fieldwork. Limited to administrative data, such as the census, it is convenient when we can justify relying on administrative boundaries. If local officials have decided to divide their city into neighborhood planning areas and to provide data at that geographic level, those areas become politically meaningful to some degree and also practical for researchers. More routinely we only have data for census tracts or similar units, and the lack of alternatives makes it legitimate to treat those tracts as if they were “real” neighborhoods.

Using census data it is easy to distinguish a subset of tracts that are candidates to be considered to be within ethnic neighborhoods for a given group. But just where should we draw the line? How Chinese or Filipino must a Chinese or Filipino neighborhood be? There is no established criterion. In their study of Italians, Irish, and Germans in the New York region, Alba, Logan and Crowder (1997, p. 892) operationalized an ethnic neighborhood as “a set of contiguous tracts, which must contain at least one tract where a group is represented as 40% or more of the residents and whose other tracts each have a level of ethnic concentration among residents of at least 35%.” However, only a handful of census tracts in Los Angeles meet this criterion for Chinese or Filipinos, because these groups are much smaller than the white ethnic groups in New York (though much of Los Angeles would be defined as a Mexican neighborhood by these criteria).

It is widely understood that the group does not necessarily have to be a majority in its identified ethnic neighborhood (a corollary is that some zones may contain “ethnic neighborhoods” of more than one group). Philpott (1978) has pointed out that the principal

Swedish ghetto identified by Park and Burgess in Chicago in 1930 was only 24% Swedish; the German ghetto was only 32% German. Some places today have international reputations as ethnic neighborhoods despite having modest proportions of group members. For example, parts of Los Angeles “are so heavily identified with Armenians that when prospective emigrants in Armenia or Iran are asked about their destination, they may answer ‘Hollywood’ or ‘Glendale,’ respectively, instead of America” (Bozorgmehr, Der-Martirosian, and Sabagh 1996, p. 368). Yet in 1990, Armenians made up only about 25% of residents of Hollywood and Glendale, reaching a maximum of 33% in their most “Armenian” tract, and only 10–15% in their peripheries. Among well-known contemporary Chinese neighborhoods, the core immigrant area of Flushing (in Queens, New York) studied by Zhou (1992) was only 14% Chinese in 1990. Monterey Park, California, was less than 25% Chinese in the mid-1980s when Horton (1995) began to study it. A recent study of minority groups in Los Angeles defined Asian residential enclaves as areas that were as little as 10% Asian (Bobo et al 2000).

Should neighborhoods be built up from tracts or smaller units within tracts? Hipp (2007) showed that, if we define neighborhoods in terms of their *effects* on residents’ experience of crime, racial/ethnic composition is more significant at the scale of the census tract than the block. But to discover the effects of average income on experience of crime, research must be done at the more localized scale of the local block. On the other hand, even the tract may prove to be too small a unit, because some social phenomena are organized at a larger scale. For this reason, recent developments in spatial statistics have encouraged the use of tract data in conjunction with information about surrounding tracts. As Sampson (1999) puts it, small areas are “embedded” in a larger urban context. That is, real neighborhoods – in the sense of the areas whose characteristics can be consequential for residents – may be larger than a single tract. Logan, Alba and Zhang (2002) found that thematic maps of New York and Los Angeles in the 1990s showed visible concentrations of several major immigrant minority groups that typically extended across many tracts. They used a measure of local spatial clustering at the tract level (local Moran’s i) to identify statistically significant clusters and treated these larger areas as the groups’ ethnic neighborhoods (see also Logan and Zhang 2004).

Geographers have also made significant advances in computational regionalization (see Guo 2008). Yet regionalization algorithms have seen limited use in the broader urban analytic community because they typically create groups out of groups. That is, because of data constraints, most of the data available to researchers are already reported over geographic regions and regionalization has not been seen as necessary. Regionalization is of most value when applied to points. Partly for this reason, when Martin (1998) sought to construct neighborhoods from census data, he began by disaggregated the reported data to raster grid cells by interpolating from enumeration district centroids. He then transformed the rasterized maps through regionalization. The same approach has been adopted by some researchers who wish to develop spatial measures of segregation (Reardon and O’Sullivan 2004).

We do not presume to know a priori the appropriate spatial scale, whether neighborhoods are likely to be found within tracts or spanning across tracts. The special advantage of analyzing point data is that we should not need to make that determination in advance. However we do make one substantive decision: the building block of neighborhoods in our research is households along a street segment. In this respect we agree with Grannis (2009, pp. 8–9) who argues that neighbor relations develop when people’s “lifestyles cause them to casually and unintentionally encounter each other and thus to have the opportunity to learn about each other through observation and to acknowledge each other’s presence or choose not to.” Neighborhoods extend beyond face blocks through a sort of ripple effect: people on one block may become acquainted with those on the next block, who are in turn tied to those

on the next, and so forth. The street network guides the paths of people's daily lives, and it is reasonable to expect neighborhoods to build up from street segments. This is not a necessary condition for any of the models that we present below; rather it is our theoretical starting point.

Our review of the literature suggests three main questions about ethnic neighborhoods. First, are their boundaries "rolling" (following Hunter, 1974) or discrete (as suggested by Suttles, 1972)? This may not be a question of how things are actually experienced on the ground but rather a question of representation: should we model ethnic neighborhoods as discrete territories or as areas with "fuzzy" boundaries? Second, how "ethnic" does a neighborhood have to be to be given an ethnic label? Our review suggests several approaches to this question, and suggests that low intensity ethnic settlement can warrant designation as an ethnic neighborhood. Third, our literature review raises questions about the scale of ethnic neighborhoods - what are the "atoms" that constitute a neighborhood- streets, buildings, people? How geographically extensive do neighborhoods tend to be? There is no "correct" answer to these questions. In the following sections we present three methods for defining ethnic neighborhoods which allow considerable flexibility in operationalization.

Data Sources

This study makes use of data assembled by the Urban Transition Historical GIS project (www.s4.brown.edu/utp). A critical component is the 100% digital transcription of records from the 1880 Census that was organized by the Church of Latter Day Saints and prepared for scholarly use by the Minnesota Population Center (MPC). This file includes approximately 50 million Americans, organized by household and with information on residents' name, age, race, gender, relation to head of household, state or country of birth, each parent's state or country of birth, and occupation, and the enumeration district, ward, city, county, and state of residence. It is available from the North Atlantic Population Project (NAPP: <http://www.nappdata.org>).

We rely on the person's and parents' race and place of birth to create categories of race and ethnicity. We create five exhaustive and mutually exclusive categories. 'Yankees' are white, born in the United States, and whose parents are also native born white. 'Irish' and 'Germans' are also white. For the foreign born, their country of birth determines their ethnicity. For those who were born in the United States but at least one parent was born abroad, the person's ethnicity is primarily determined by their mother's country of birth. If only the father was foreign-born (or if the mother was foreign-born but her birthplace was not reported), the father's country of birth will be applied. People of other than the above three race/ethnic groups are categorized as a single residual group of 'all others'.

The 1880 census did not gather information on income or education, the most conventional indicators of socioeconomic standing. The available measure is the socioeconomic standing of the person's occupation on a scale of 0–100occupational (SEI) provided by MPC, based on the average education and earnings of persons in each occupation as measured in 1950. An important question in using the SEI in a historical study is whether the relative standing of occupations is stable over time. Sobek (1996) has studied this question directly, comparing the average income of men in each of 140 occupations in 1890 to the income of men in those occupations in 1950. The correlation between the two is .93.

Addresses for residents of Newark (and 38 other cities not used here) were transcribed and merged with the information about residents in the NAPP files. We juxtaposed a map showing 1880 Essex County boundaries with the US Census Bureau TIGER street files (<http://www.esri.com/data/download/census2000-tigerline/index.html>). The street file was edited by hand to reconstruct the street layout of Newark in 1880. Doing this involved

deleting streets or highways constructed after 1880, adding those that were demolished, changing names of others, and in some cases correcting the alignment of streets. The Newark city directory for 1880 included a listing of address ranges by block for most streets, and this information was used to geocode addresses. Additional checks were made to ascertain the boundaries of enumeration districts and ensure that individuals were placed in their recorded enumeration district.

Figure 1, a map of several streets in Newark, illustrates the resulting data. For any single building we know its location and the characteristics of every resident, including what household they live in (there are clearly at least three households at the selected address in the figure, 41 Walnut Street). Key characteristics used here are age, ethnicity, and SEI. In this figure we have labeled buildings and street segments according to the dominant ethnicity of residents. At the selected address 8 of 10 adult residents are Yankee. Four residents, all women, have no occupation (SEI=0). Some others (Fred and Gabriel Thorn) have very high status occupations, like doctors or judges.

Table 1 reports summary information on the 1880 population of Newark. The total city population was 133,554, of whom 80,116 were adults (age 18 and above). Residents were fairly evenly divided among the three main ethnicities: Germans were the largest group, followed closely by Yankees (native born whites with native parents), and Irish. About 20% of the population was classified into other races and ethnicities, of which the largest (about 10%) was British. For the sake of simplicity, we entirely disregard this other population in the following analyses. The following analyses are based on the adult population (age 18 and above), whose ethnic composition is very similar to that of the total population.

In parts of our analyses we examine the ethnic composition of street segments in Newark. Newark had 1780 street segments with at least one Yankee, German or Irish adult resident. These have been classified in two different ways. In the energy minimization model each street segment begins with a “label,” and for our purpose the label denotes the dominant resident group. Street segments composed 50% or more of one group have been labeled as Yankee, German or Irish; all other street segments have been labeled as mixed. For the Bayesian model there are three independent categorizations of street segments, and these reflect the clustering of members of each ethnic group. For example, a street segment with more than a given number of Irish residents is classed as “high” Irish, regardless of how many residents represent other groups (other categories are medium and low). That street segment has a larger than average number of Irish, but it is not necessarily predominantly Irish.

The unique disaggregated nature of the Urban Transition HGIS data frees us from thinking about neighborhoods as territorial units defined by geographic boundaries. Instead we can define neighborhoods based upon the ethnic composition of an area, potentially on a building-by-building basis. We define our problem as how to combine buildings or streets into some larger unit while maintaining flexibility in terms of scale, the nature of neighborhood boundaries, and the number and composition of neighborhoods.

Studying neighborhoods through k-functions

A single building or street does not by most definitions constitute a neighborhood. Our first approach to the creation of neighborhoods is rooted in “local k-functions” (Franklin and Getis 1987). A k-function measures the expected number of “events” around an “arbitrary event” as a function of distance. The term “event” refers to occurrences that have a probability of occurring at a geographic location. K-functions are typically used to understand the geographic pattern of such events as disease, accidents, or crime. The k-function is estimated by counting up the number of events around each event and dividing

by an intensity parameter λ , which measures the density of events in the study area (Bailey and Gatrell, 1995). K-functions are typically used to answer questions about the geographic distribution of events, especially whether they are clustered or occur randomly in space.

In our case we are not working with events like muggings or car accidents. Our events are buildings, and their location is determined by exogenous factors like the layout of streets. We use the k-function in a somewhat atypical way, rather than looking for places where we see more/less events than expected we use the k-function to describe the character of the neighborhood around each building. We compute thousands of *local* k-functions to provide a statistical summary of the prevalence of the Irish, German, and Yankee ethnic groups around every location in the study area (by contrast to a global k-function which would summarize an entire study area). More concretely we draw a series of concentric rings around each of the 15000+ residential buildings in Newark. The rings count up the number of people by ethnicity within 50, 75, 100, 200, 300, 400, and 500 meter radial buffers. The result is a series of three graphs (one for each ethnicity) for each building. To create realistic graphs we had to consider the scale of the environment experienced by a resident of Newark in 1880. Hershberg (1979, p. 136) estimated the average commute in Philadelphia in for different professions in 1880 to be between .2 km (for blacksmiths) to 1 km (for lawyers), and we adopted the 1km diameter for our largest distance band. This means that we are including in each building's "catchment area" the neighborhood scale that a typical working adult might experience in a given day. Within this catchment area we measure ethnic prevalence at regular intervals, oversampling at the sub-block (< 100m) scale.

Our k-functions include a constant intensity parameter (λ). This means that in densely settled parts of the city the observed values of the k-function will tend to be higher simply because the total population density is higher. In addition the observed values for points near the edge of the study area will be lower because their concentric rings include areas that were not observed. We control for both of these distortions by calculating two different k-functions: an ethnic-specific one and another for the total population, as shown in equations (1) and (2) below.

$$K_{ij}(h) = \frac{1}{N_i} \lambda^{-1} (\# \text{ of people of ethnicity } i) \quad (1)$$

$$K_{tj}(h) = \frac{1}{N} \lambda^{-1} (\# \text{ of people}) \quad (2)$$

Where K_{ij} is the k-function for ethnicity i at location j and K_j is the k-function for the total population at location j . N is the total population and N_i is the total population of ethnicity i . It is the difference between these k-functions that we use in our analysis. Conceptually, this is similar to the $L(h)$, an adaptation of the k-function whereby the $E\{k(h)\}$ under complete spatial randomness is subtracted from the observed value of $k(h)$ (Bailey and Gatrell, 1995). In addition we normalize the densities so that each distance band has a zero mean and unit variance.

For example, Figure 2 shows the k-functions for 41 Walnut Street (the selected location in Figure 1). As previously noted, this building itself has a high prevalence of Yankees at all distance bands, and Figure 1 seemed to show many Yankee buildings around it. Figure 2 quantifies this observation with 3 k-functions, one for density of Yankees, another for density of Irish, and a third for density of Germans. It shows that the line representing the Yankee k-function has values above 0, and the density rises rapidly to 100 meters. In

contrast the lines are below 0 for Irish and German ethnicity. With only this information, one might guess that this building is in a “Yankee neighborhood.”

However Figure 2 displays only 3 of the roughly 45,000 local k-functions we computed to describe Newark. Figures 3–5 show similar graphs for three other buildings. Figure 3 shows a house at 39 Belmont St. that has a high prevalence of Germans at all distance bands, especially at 300 meters. Figure 4 (144 Newark Street) shows a case where the highest prevalence is Irish, especially at above 75 meters’ distance. And Figure 5 (83 Chambers street) show a more mixed case where the highest density of neighbors is German up to about 300 meters, but beyond that distance the densities of Germans and Yankees are equal.

By design local k-functions have a high degree of spatial autocorrelation; that is, those at locations near each other tend to be similar. This is a very desirable property for our purpose because it means we can expect to find spatial clusters of buildings with similar k-functions and use these clusters to identify types of neighborhoods. If a city has ethnic districts we would expect the k-functions to reveal regions where the curve for one ethnicity is dominant. If however, the city is spatially undifferentiated we would not expect clearly defined regions where one ethnic curve is measurably higher than others. Cases like Figure 5 would be more common than cases like Figures 2–4. Burstein (1979) notes that 1850–1880 was a time when American cities were in transition from a nineteenth century form that was not residentially differentiated by ethnicity or class toward a twentieth century pattern with clearly defined class and ethnic zones. Were there clearly defined ethnic neighborhoods in Newark, NJ in 1880?

Let us think of ethnic neighborhoods as zones of the city where a specific pattern of k-functions tends to cluster. We make no a priori assumptions about what are the types of patterns, but we do set the number of patterns that should be distinguished. In this example the number is four, and we are looking for Yankee, German, Irish, and mixed neighborhoods without defining them in advance. A cluster is defined simply as buildings with “similar” k-functions. The solution is provided by a simple k-means cluster analysis on the 45,000 local k-functions.

K-means is a way to partition a data set into a user specified number of groups, where the groups are created through an algorithm that assigns observations to a groups based on how similar they are to the mean of each group (Romesburg, 2004). The procedure is iterative, after the initial assignment of observations to classes the mean of each group is recalculated and observations are re-assigned. The procedure terminates once the classes become stable and there is little or no reassignment of observations between iterations.

The input to the k-means cluster analysis was the observed value of the local k-functions for each of three ethnicities at each of seven distance bands for all 15,000 buildings in Newark. The groups that emerged from the k-means cluster represent observations that have similar characteristics. Mapping the 4-cluster result (Figure 6) shows 4 clearly defined regions; each of these regions corresponds to groups of buildings with similar k-functions.

This map paints a picture of the ethnic spatial structure of the city. A section of the city running north-south and not far from the river on the east end is identified as Yankee. This identification is based on inspecting the typical k-function curves of buildings in this zone such as 41 Walnut Street. For example, we could easily plot the average densities at every distance band for all buildings in this zone. There is one large concentrated area in the western part of the city identified as German (and this is where 39 Belmont Avenue is found). There are four or five areas identified as Irish, including the area where 144 Newark Street is located. And finally there are other areas designated as mixed, including the vicinity of 83 Chambers Street.

This use of local k-functions offers one answer to the question of how Newark neighborhoods were structured by ethnic composition in 1880. It identifies areas that can be treated as “ethnic neighborhoods” and provides a guide to their boundaries. The extensive territories that are categorized as “mixed” suggest that there are indeed intermediate areas, or perhaps zones of transition, around the borders of these neighborhoods, which is consistent with urban theory. If desired, it is possible to search for more complicated patterns – for example, to see whether allowing a larger number of neighborhood types would produce clusters that reflect Irish-Yankee or Irish-German combinations. It is also possible to extend the analysis to include not only ethnicity but also other dimensions such as the distinction between first and second generation, or gradations of occupational SEI.

The k-functions approach is superior to a simple mapping of the ethnicity of buildings, street segments, or larger areas on the street grid, because it takes into account the full array of other street segments at varying distances. Because spatial dependence is built into the k-functions for street segments that are near to each other, the method has a tendency to smooth out local variations and yield the simpler picture that social scientists have in mind when they speak of ethnic neighborhoods.

Energy minimization

We turn now to a second approach that uses the same input data to produce a somewhat different map of ethnic neighborhoods. Energy minimization refers to an optimization technique that is most commonly used to process digital images such as medical images (Boykov et al 2001). Its assumption is that a pixel that is unlike adjacent pixels may represent the edge of an object (like a tumor) or could just be noise (measurement error). Information from surrounding pixels is used to refine images, pixels are reclassified, or in the terminology of this method, relabeled (in our application, a label describes the ethnicity of a neighborhood).

As applied to our research, we do not literally believe that the ethnic composition of a street segment is measurement error. Rather, as an indicator of whether the location falls within a largely Yankee, Irish, German, or mixed neighborhood, the information is incomplete. On the one hand, an all-Irish street segment has some probability of signifying an Irish neighborhood. On the other hand, the characteristics of adjacent street segments also count. The question is how to balance these two criteria.

In formal terms, we determine if a street is part of a larger neighborhood by minimizing the following quantity:

$$E(f) = E_{smooth}(f) + E_{data}(f) \quad (3)$$

Here $E(f)$ is the net energy cost of the neighborhood label that is assigned to a street segment. This cost has two components. The first, $E_{smooth}(f)$, controls the smoothness of the resulting image (or map). A smooth map consists of large homogenous regions whereas a rough map has a patchwork of small neighborhoods. The smoothness term applies a penalty for adjacent street segments that are labeled differently but no energy (cost) is added if they are labeled the same. For a given street the smoothness term considers all connecting street segments, and the costs associated with every one are part of the net energy. We are assuming that cities consist of ethnically homogeneous neighborhoods and so we assign a “cost” to calling a street “Irish” if it is surrounded by streets where German is the dominant ethnicity.

The second component, $E_{data}(f)$, is referred to as the data cost. It assesses a penalty for assigning to a street a label that is different from the observed value. That is, there is a cost of labeling a predominantly Irish street segment as being in a German neighborhood. The notion is that the observed values have significance. The sum of the data and smoothing costs for each street is the quantity that we wish to minimize, and that controls the final map that the method produces. Like k-functions, it is up to the researcher to decide which street segments are near enough to consider as relevant neighbors; in this example we only count street segments that are connected at an intersection. The researcher also controls the parameters that are used to calculate each cost. The final map depends on how one chooses to weigh the smoothing costs vs. the data costs.

The approach is illustrated in Figure 7. Here Street A is observed from census data to be predominantly Irish (i.e., more than 50% of adult residents are Irish). However three of the four connecting streets are German. We ask, “Is Street A in an Irish or a German neighborhood?” The figure shows two scenarios in which the relative weights of smoothing cost and data cost are different. Under scenario 1 the cost of changing the label for Street A (from the observed Irish to a potential German) is set at 20 and the cost of being connected to each street that is a different ethnicity is 4. The total energy cost is easily calculated to be 12 if Street A is labeled Irish, but 24 if it is labeled German. Using these parameters, Street A remains Irish. Under scenario 2 there is a higher cost for having different types of streets connected to each other (12) but the same data cost (20). In this scenario the total energy cost would be minimized by labeling the street German.

We use the Boykov et al. (2001) -expansion algorithm to assign ethnic labels to streets that minimize the “energy” (equation 3) of the resulting map. The input data for this analysis is shown in Figure 8, and the result is mapped in Figure 9. The result is a product of some assumptions. Our initial classifications were based on the observed proportion of each ethnic group on the 1780 populated street segments of Newark, classifying each street segment as Yankee, German, Irish (50% or more of adults were of that ethnicity) or mixed. We experimented with many combinations of smoothing and data costs. The outcome shown here strikes what we assess to be an appropriate balance between a rough and smooth map.

Note that compared to the input map, the result of energy minimization is a simplified depiction of the ethnicity of different sections of Newark. Many street segments whose actual composition was unlike adjacent segments have been relabeled, and a large zone has been labeled as mixed, even though only a few street segments began in the mixed category given how we defined that category. Very few street segments appear as isolated enclaves, and these are found only on the edges of the city.

The Boykov -expansion algorithm produced many maps that also minimized Equation 3 (using different sets of parameters) but did not correspond to our understanding of the city based on detailed inspection of the input data. The connectivity of the street network in a given city, the number and definition of categories used for the input data, and the relative difference between the data and smoothing costs all affect the resulting map. In this sense Figure 9 is arbitrary. However it is not necessarily more arbitrary than other methods, and the iterative process of varying and evaluating the result offered a systematic way of probing the spatial pattern of Newark. The figure shows clearly defined Irish, Yankee, German, and mixed neighborhoods rather similar to the local k-function result, which is also reassuring.

Both of these approaches, local k-functions and energy minimization, give different results depending on how they are applied. It would be preferable to find historical, empirical, or theoretical justification for a given set of parameters and assumptions. In image processing it is often possible to objectively compare the result of an algorithm to “ground truth” because

one knows more or less what the picture should look like. A better grounded approach to identifying ethnic neighborhoods with these methods would be possible if we had additional information to validate the results, such as archival sources giving the location of ethnic institutions like churches or stores that social scientists understand to be additional markers of ethnic turf. We return to this possibility in the conclusion.

Bayesian model of neighborhood clusters

The third and final approach is based in Bayesian analysis, which has been widely applied in spatial statistics (Congdon 2003). Bayesian analysis is about making statistical inference from combining researchers' prior beliefs and sample data (Gelman et al. 2004). Its crucial feature is that any unknown parameter in a Bayesian model is considered to be a random variable. Any information about the parameters before collecting the data such as previous findings and expert knowledge is specified as a "prior probability distribution." The process of inference consists of using the sample data to adjust prior beliefs and deriving the "posterior distribution" of the parameters. We draw here on an application of Bayesian methods by Xu and Bao (2004), who used this approach to detect spatial clusters and identify disability patterns at the county level in Mississippi and Alabama.

The basic unit of analysis here is the street segment. As in the energy minimization example above, we define the neighboring street segments as those which are directly connected to the focal street segment, and the ethnic classification of a street segment in the observed data is based on the ethnicity of the adult residents.

The input data are of the following form: among n_i people living along street segment i , there are y_i people of a given ethnic group. We posit that the probability for observing this composition is governed by the underlying spatial pattern of ethnic neighborhoods. That is, we posit a "real" ethnicity of neighborhoods and treat the observed values on a given street segment as its realization with some random variation. More specifically our approach here is to think of neighborhoods' ethnicity in three dimensions, its degree of Yankee-ness, German-ness, and Irish-ness. On each dimension neighborhoods are assumed to be high, medium or low (i.e., ethnic, mixed, or non-ethnic). This means, for example, that in theory a neighborhood could be high ethnic Yankee and high ethnic German at the same time. For each group let p_0 , p_1 , and p_2 denote the probabilities of observing y_i group members living among a total number of n_i persons in each of these three types of neighborhoods.

These probabilities must be estimated separately for each group. Under a Bayesian framework, any unknown quantity is assumed to be random. (In our application we make another assumption, that for a given group the probabilities are constant throughout the city.) Our task is to estimate the distribution of this random quantity and draw summary inferences from this distribution for each street segment. This is a complex task, but fortunately it can be accomplished using established procedures, as follows.

First, following Xu and Bao (2004), we consider the following binomial distribution for the number of people from an ethnic group given the underlying k types of neighborhoods:

$$P(Y_i=y_i|X_i=k)=\binom{n_i}{y_i}p_k^{y_i}(1-p_k)^{n_i-y_i} \quad (4)$$

X denotes the random set of ethnic neighborhood configurations. Given the underlying process for neighborhoods $X = x$ in one realization, we assume that X takes on values 0, 1, and 2 for the three types of neighborhood (k). Equation (4) shows that we estimate a binomial model for the probability of observing a given number of ethnic individuals in each

street segment. If we ignored the spatial structure of the data, we would simply estimate a conventional binomial model as for other dichotomous outcomes in social science research. We would be using aggregate data at the street segment level, similar to the way one models binary outcomes from a contingency table (Agresti 2002).

Second, we incorporate the spatial structure by imposing a neighboring effect in the model. A Potts model is specified as the prior distribution for the underlying neighborhoods X . The Potts model has the property of a Markov random field (MRF) which simply means that streets that are near each other are more likely to be assigned the same type of neighborhood than distant streets (François et al. 2006: 806). The MRF property makes the Potts model a good choice for modeling spatial data (Cressie 1991, Xu and Bao 2004; François et al. 2006). If x is one possible realization of the random set X with each component being one of the k three types of neighborhood, then the prior distribution is given by

$$P(X=x)=Z(\rho)e^{\rho\sum_{k=0}^2\sum_{i\sim j}I_k(i)I_k(j)} \quad (5)$$

Here ρ is a parameter that specifies strength of the spatial autocorrelation between neighboring areas. $Z(\rho)$ is a normalizing constant that can be dropped in our Bayesian model specification and estimation (Gelman et al. 2004). Neighborhood relationships are captured through a contiguity based indicator function, $I_k(i)$ is the indicator function which is 1 if $x_i = k$ and 0 otherwise. Here, $i \sim j$ means that in this case, street segments i and j are neighbors, provided that they are directly connected to each other through an intersection. We consider global, instead of local, spatial autocorrelation for each ethnicity in the study area (Liu 2001) and use Moran's I computed using the first order contiguity street segments as an estimate for ρ (Xu and Bao 2004).

Third, a useful feature of this model is that it facilitates incorporating covariates into the model by linking $\text{logit}(p_i)$ or $\text{probit}(p_i)$ to a linear combination of predictors Z . Model estimation is done by using Gibbs sampling, a particular useful Markov chain Monte Carlo (MCMC) algorithm for solving multidimensional problems (Geman and Geman 1984). We will illustrate the introduction of covariates by controlling for SEI of the street segment (the mean for employed adults) in the final step of the analysis.

Figure 10 depicts the results for an area of Newark in 1880. Map A shows the initial Yankee categories for these street sections, divided into three equal terciles based on the number of Yankees (low, medium and high). The same procedure is performed for Germans and Yankees separately.¹ Note that street segments with a high number of Yankees may not have a high proportion of Yankees if the density of Irish or Germans is also high. Map B shows the Bayesian solution (the posterior median estimates of classifications of street segments for Yankees) without controlling for SEI. It is notable that quite a few street segments initially classified as having high proportion of Yankees in Map A have been relabeled in this solution as having medium or low proportion, and some street segments with medium proportion of Yankees have been relabeled as low density. This suggests that some street segments may have been “mistakenly” classified in the input data, provided that a binomial model is appropriate.

¹For Irish, a street segment is initially classified as low if it has 2 or fewer Irish, medium if it has 3–9, and high if it has 10 or more. For Germans, a street segment is initially classified as low if it has 2 or fewer Germans, medium if it has 3–11, and high if it has 12 or more. For Yankees, a street segment is initially classified as low if it has 2 or fewer Yankees, medium if it has 3–13, and high if it has 14 or more.

Map C shows the posterior median of the types of street segments for Yankees after controlling for the average SEI of street segments in the same area as. It is notable that some street segments that are classified as having high (or medium) proportion of Yankees in Map B now are labeled as medium (or low) proportion. This suggests that for some street segments, the spatial concentration of Yankees can be attributed to the concentration of individuals of similar socioeconomic status (in particular, high status). In other words, Yankees may live close to each other not because of their preference for the same ethnicity but because of their preference for the same socioeconomic status group. Given a certain level of SEI (i.e. controlling for SEI), the proportion of Yankees is no longer great enough for the street segments to be classified as high.

Table 2 summarizes the classifications of street segments in the initial Bayesian solution and in the solution controlling for SEI. It shows that about 4% of street segments have been relabeled from having medium to low, and 6.6% from having high to medium proportion of Yankees after controlling for SEI. On the other hand, about 8.7% of street segments for Irish have been reclassified from low to medium, and 4.7% from medium to high after controlling for SEI. This suggests that for some street segments, the spatial concentration of Irish can be confounded by the concentration of similar socioeconomic status. Some street segments actually have relatively high proportions of Irish, but they could be mistakenly classified as having low proportions because they appear to be clusters of low SEI individuals who are also Irish. Such flips (only about 3%) rarely happen to Germans. This is consistent with the fact that Yankees were on average of the highest socioeconomic status, Irish had the lowest status, and Germans ranked in the middle in 1880 Newark. The fact that only a relatively low proportion of street segments become reclassified after controlling for SEI suggests that SEI matters much less than ethnicity itself.

We create city-wide maps of Irish-ness, German-ness, Yankee-ness. We can then overlay these maps to identify whether an area is “mono-ethnic” (high on only one dimension) or mixed. We finally overlay the three maps of posterior median of ethnic proportions for Irish, Germans, and Yankees after controlling for SEI on top of each other. We can classify the street segments that are labeled as having high proportion of Irish and low proportion of both Germans and Yankees as “Irish” neighborhoods, and so on for “German” and “Yankee” neighborhoods, resulting in a map as in Figure 11. “Other” neighborhoods are those in which the density of no group is high. This specific solution could yield different maps. For example, we could add additional categories such as high on two or three ethnicities to highlight particular combinations of groups. The overall pattern is similar to those obtained from K-functions and energy minimization, although the resulting ethnic neighborhoods seem to be more fragmented, especially for some Irish and Yankee street segments.

Discussion and conclusion

Geographic information technologies have led to an explosion of new data sources, allowing us to re-ask old questions about the organization of cities. The study of the ethnic differentiation of cities into neighborhoods has heretofore largely depended on geographic aggregations of people. Whether these aggregations are census tracts or wards or some other unit, they become the lens through which the city is viewed. They constrain what we can see because we cannot look inside them. They also limit how we can think about neighborhoods. Many social scientists begin with a concept of neighboring: the connections and daily interactions between households living in the same building, or the buildings next door or across the street, or maybe in the next block. But our data impose a different notion that is well captured in the nature of the “exposure indices” that are a standard measure of segregation. A higher exposure index means that members of one group tend to live in areas

with larger shares of another group. But we know this doesn't mean that there is actually more social interaction between these groups, especially because the geographic scale at which the index is calculated may not be related to the scale of neighboring.

Census data will never directly measure social interaction. But the increasing availability of fine-grained information on individual people and households offers the possibility of constructing maps of neighborhoods that begin at a more natural scale. At the same time, having so much information also puts more pressure on our ability to use it. Dealing with 100 or 500 census tracts in a city is much simpler than dealing with 100,000 locations. This is the motivation for our experimentation with alternative methods of neighborhood identification, to learn how to study spatial patterns at a much finer level while linking operational decisions to theoretical conceptions of what a neighborhood is.

We used three different methods to reveal the underlying social structure of the city. The Bayesian and energy minimization approaches assume that the types of neighborhoods that exist within a city are known *a priori*. The question they seek to answer is, "where are the ethnic neighborhoods?" The k-function approach, on the other hand, does not make assumptions about the number and/or types of neighborhoods – it asks, "what are the neighborhoods and where are they located?" It is not our intention to argue for the superiority of any one of these approaches. They all begin with the same information about Newark and it is not surprising that the resulting maps are generally similar.

Since there is no conventional template for assessing these models, the right result could be seen as indeterminate. There are potential empirical guides to making the "right" choice. For example, as in many cluster analyses one might use the criterion that the result should maximize variation between neighborhoods while also maximizing homogeneity within them. However this criterion presumes that each neighborhood should be composed mainly of one dominant ethnic group, while in principle it is possible that there are socially significant transitional zones or other sorts of mixed neighborhoods. Another more substantively appealing criterion would employ external validation. Suppose, for example, that archival research could determine the locations of churches whose denomination or name identifies an ethnic affiliation, or the addresses of members of ethnic voluntary associations. One criterion could be that the best mapping of neighborhoods is the one that locates these institutions most centrally in zones represented by their ethnic constituents. Historical documents that give ethnic labels to areas of the city would also be of use for validation.

Ultimately, though, empirical indicators (like the labels that our methods apply to buildings or street segments) are best judged by their consistency with a theoretical concept. Let us review the assumptions made by each method and its operationalization here.

The first question raised in the introduction is whether neighborhood boundaries are rolling or discrete. The Bayesian method views neighborhoods probabilistically: a street segment is not in or out of a neighborhood but rather has a probability of belonging to one of 9 neighborhood types (high, medium, low intensity of each ethnicity). This method suggests that a neighborhood is not a discrete territory, but rather a field of varying membership probabilities.. Zones where these probabilities are low could be understood as zones of transition or as borders. By contrast, in the energy minimization and k-function approaches neighborhood assignment is discrete not probabilistic.

The second question is how ethnic a neighborhood needs to be in order to be given an ethnic label. An advantage of all of the methods presented here is that they do not require setting an *a priori* threshold of ethnicity. Instead the question is framed in terms of some form of

spatial autocorrelation: from knowledge about “who are your neighbors” we infer what the neighborhood is.

The third question is what is the scale of the neighborhood and what are its components. Another advantage of all three approaches here (and of the underlying dataset) is that geographic scale is not predetermined, and any of the approaches could in principle be applied from information about buildings or street segments, or any other geographic unit. In one case we began with data about individual buildings and asked about their relation to other buildings at a variety of spatial scales. In two other cases we presumed that nearby buildings should be aggregated into street segments, that it is along the two sides of the street on the same block that people are likely to have the most face to face interaction and therefore the same shared sense of neighborhood. We treat intersections of streets as effective delimiters of shared social space, but we also treat them as connectors. In both the energy minimization and Bayesian models we choose to treat only connected street segments as meaningful neighbors of the focal street segment. We could alternatively have stretched to further distances, taking into account the ethnic composition of the next connected street segment, and the next. In the k-function analysis we created distance bands as concentric rings, but measuring those distances along the street network might have better represented our expectations about how streets shape interactions. With any of these approaches we could have adopted the view that “neighbors” at all distances count but that nearer ones count more.

We present these alternatives and the choices we made in operationalizing them as a step toward better understanding of neighborhoods, emphasizing that we do not believe there is one right answer. The principal challenge for researchers interested in using large high resolution geographic data bases to study cities is connecting theory to methods and data. Selecting an approach and applying it to data is a thoughtful process, and the right one is the one that most clearly reflects the way the analyst thinks about underlying social processes.

Acknowledgments

This research was supported by research grants from National Science Foundation (0647584) and National Institutes of Health (1R01HD049493-01A2) and by the staff of the research initiative on Spatial Structures in the Social Sciences at Brown University.

References

- Agresti, Alan. *Categorical Data Analysis*. 2. Hoboken, NJ: John Wiley & Sons, Inc; 2002.
- Alba, Richard D.; Logan, John R.; Crowder, Kyle. White Neighborhoods and Assimilation: The Greater New York Region, 1980–1990. *Social Forces*. 1997; 75:883–909.
- Bailey, Trevor; Gatrell, Anthony. *Interactive Spatial Data Analysis*. Englewood Cliffs, NJ: Prentice Hall; 1995.
- Bobo, Lawrence D.; Oliver, Melvin L.; Johnson, James H., Jr; Valenzuela, Abel. *Prismatic Metropolis: Inequality in Los Angeles*. New York: Russell Sage Foundation; 2000.
- Boykov, Yuri; Veksler, Olga; Zabih, Ramin. Fast Approximate Energy Minimization via Graph Cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2001; 23:1222–1239.
- Bozorgmehr, Mehdi; Der-Martirosian, Claudia; Sabagh, George. Middle Easterners: A New Kind of Immigrant. In: Waldinger, Roger; Bozorgmehr, Mehdi, editors. *Ethnic Los Angeles*. New York: Russell Sage Foundation; 1996. p. 345–378.
- Burstein, Alan N. Immigrants and Residential Mobility: The Irish and Germans in Philadelphia 1850–1880. In: Hershberg, Theodore, editor. *Philadelphia: Work, Space, Family, and Group Experience in the Nineteenth Century*. New York: Oxford University Press; 1981.

- Campbell, Elizabeth; Henly, Julia R.; Elliott, Delbert S.; Irwin, Katherine. Subjective Constructions of Neighborhood Boundaries: Lessons from a Qualitative Study of Four Neighborhoods. *Journal of Urban Affairs*. 2009; 31:461–490.
- Congdon, Peter. *Applied Bayesian Modelling*. New York: Wiley; 2003.
- Cressie, Noel. *Statistics for Spatial Data*. New York: Wiley; 1991.
- Fotheringham AS, Wong DWS. The modifiable areal unit problem in multivariate statistical analysis. *Environment and Planning A*. 1991; 23(7):1025–1044.
- François, Olivier; Ancelet, Sophie; Guillot, Gilles. Bayesian clustering using hidden Markov random fields in spatial population genetics. *Genetics*. 2006; 174:805–816. [PubMed: 16888334]
- Gelman, Andrew; Carlin, John B.; Stern, Hal S.; Rubin, Donald B. *Bayesian Data Analysis*. 2. New York: Chapman Hall; 2004.
- Geman, Stuart; Geman, Donald. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1984; 6(6):721–741. [PubMed: 22499653]
- Getis A, Franklin J. Second-Order Neighborhood Analysis of Mapped Point Patterns. *Ecology*. 1987; 68(3):473–477.
- Grannis, Rick. *From the Ground Up: How the Layered Stages of Neighbor Networks Translate Geography into Neighborhood Effects*. Princeton: Princeton University Press; 2009.
- Guo D. Regionalization with dynamically constrained agglomerative clustering and partitioning (REDCAP). *International Journal of Geographical Information Science*. 2008; 22(7):801.
- Hartshorne, Richard. *The Nature of Geography; a Critical Survey of Current Thought in the Light of the Past*. Lancaster, PA: The Association of American Geographers; 1939.
- Hershberg, Theodore; Light, Dale; Cox, Harold E.; Greenfield, Richard. The 'Journey-to-Work': An Empirical Investigation of Work, Residence, and Transportation, Philadelphia, 1850 and 1880. In: Hershberg, Theodore, editor. *Philadelphia: Work, Space, Family, and Group Experience in the Nineteenth Century*. New York: Oxford University Press; 1981.
- Hipp, John R. Block, Tract, and Levels of Aggregation: Neighborhood Structure and Crime and Disorder as a Case in Point. *American Sociological Review*. 2007; 72:659–680.
- Horton, John. *The Politics of Diversity: Immigration, Resistance, and Change in Monterey Park, California*. Philadelphia: Temple University Press; 1995.
- Hunter, Albert. *Symbolic Communities*. Chicago: University of Chicago Press; 1974.
- King, Gary. *A Solution to the Ecological Inference Problem*. Princeton, NJ: Princeton University Press; 1997.
- Lacy, Karyn R. *Blue-Chip Black: Race, Class, and Status in the New Black Middle Class*. Berkeley, CA: University of California Press; 2007.
- Liu, Jun S. *Monte Carlo Strategies in Scientific Computing*. New York, NY: Springer; 2001.
- Logan, John R.; Zhang, Wenquan. Identifying Ethnic Neighborhoods with Census Data: Group Concentration and Spatial Clustering. In: Goodchild, Michael; Janelle, Donald, editors. *Spatially Integrated Social Science*. New York: Oxford University Press; 2004. p. 113-126.
- Logan, John R.; Zhang, Wenquan. Identifying ethnic neighborhoods with Census data. In: Goodchild, Michael F., editor. *Spatially Integrated Social Science*. New York: Oxford University Press; 2004. p. 113-126.
- Logan, John R.; Alba, Richard D.; Zhang, Wenquan. Immigrant Enclaves and Ethnic Communities in New York and Los Angeles. *American Sociological Review*. 2002; 67:299–322.
- Logan, John R.; Alba, Richard D.; Zhang, Wenquan. Immigrant enclaves and ethnic communities in New York and Los Angeles. *American Sociological Review*. 2002; 67:299–322.
- Martin, David. Automatic neighbourhood identification from population surfaces. *Computers, Environment and Urban Systems*. 1998 Mar 1; 22(2):107–120.
- Omer, Itzhak; Or, Udi. Distributive Environmental Justice in the City: Differential Access in Two Mixed Israeli Cities. *Tijdschrift voor Economische en Sociale Geografie*. 2005; 96:433–443.
- Openshaw S. A Geographical Solution to Scale and Aggregation Problems in Region-Building, Partitioning and Spatial Modeling. *Transactions of the Institute of British Geographers*. 1977; 2(4): 459–472. New Series.

- Philpott, Thomas Lee. *The Slum and the Ghetto: Neighborhood Deterioration and Middle-Class Reform, Chicago, 1880–1930*. New York: Oxford University Press; 1978.
- Reardon SF, O'Sullivan D. Measures of spatial segregation. *Sociological Methodology*. 2004; 34:121–162.
- Rich, Meghan Ashlin. 'It Depends on How You Define Integrated': Neighborhood Boundaries and Racial Integration in a Baltimore Neighborhood. *Sociological Forum*. 2009; 24:828–853.
- Romesburg, H Charles. *Cluster Analysis For Researchers*. Wadsworth; 2004.
- Sampson, Robert; Morenoff, Jeffrey; Earls, Felton. Beyond Social Capital: Spatial Dynamics of Collective Efficacy for Children. *American Sociological Review*. 1999; 64:633–660.
- Sobek, Matthew. Work, Status, and Income: Men in the American Occupational Structure since the Late Nineteenth Century. *Social Science History*. 1996; 20:169–207.
- Stoneall, Linda. Cognitive Mapping: Gender Differences in the Perception of Community. *Sociological Inquiry*. 1981; 51:121–127.
- Suttles, Gerald. *The Social Construction of Communities*. Chicago: University of Chicago Press; 1972.
- Xu, Li-An; Bao, Shuming. Detecting spatial clusters: a Bayesian approach with application to identifying disability patterns in Mississippi and Alabama. In: Getis, Arthur; Mur, Jesús; Zoller, Henry G., editors. *Spatial Econometrics and Spatial Statistics*. New York: Palgrave Macmillan; 2004. p. 265-295.
- Zhou, Min. *Chinatown: The Socioeconomic Potential of an Urban Enclave*. Philadelphia: Temple University Press; 1992.

Figure 1.
Geocoded information for a section of Newark, NJ in 1880

Figure 2.
K-functions for a Yankee building in Newark

Figure 3.
K-functions for a Yankee building in Newark

Figure 4.
K-functions for a German building in Newark

Figure 5.
K-functions for an Irish building in Newark

Figure 6.
Ethnic neighborhoods identified through local k-functions

Figure 7.
Illustration of cost calculations in energy minimization

Figure 8.
Map of input data for energy minimization

Figure 9.
Map of results from energy minimization

Figure 10.
Initial values and posterior medians of types of street segments for Yankees in a section of Newark, 1880

Figure 11.
Combined Bayesian classification of street segments for Irish, Germans, and Yankees in Newark 1880

Table 1

Ethnic composition of population and street segments, Newark 1880

	Population	Population age 18+	Classified street segments
Yankee	37,180	27.8%	22,608
German	40,042	30.0%	22,561
Irish	30,158	22.6%	18,740
Other/mixed	26,174	19.6%	16,207
Total	133,554	100.0%	80,116
			1,780
			100.0%

Table 2

Percentages of types of street segments before and after controlling for SEI

	No control for SEI			
Control for SEI	Low	Medium	High	Row Sum
<hr/>				
Irish				
Low	42.2	1.1	0.2	43.5
Medium	8.7	23.6	2.4	34.7
High	0.1	4.7	17	21.8
Column Sum	51	29.4	19.6	100
German				
Low	49.1	1.1	0	50.2
Medium	0.7	26.3	0.7	27.8
High	0	0.5	21.6	22.1
Column Sum	49.8	27.9	22.3	100
Yankee				
Low	36.4	4	0.3	40.7
Medium	3.4	18.4	6.6	28.4
High	0	3.1	27.7	30.8
Column Sum	39.8	25.6	34.6	100