Published in final edited form as:

Cartogr Geogr Inf Sci. 2010 July 1; 37(3): 169–187. doi:10.1559/152304010792194930.

Bicomponent Trend Maps: A Multivariate Approach to Visualizing Geographic Time Series

Jonathan P. Schroeder

Minnesota Population Center, University of Minnesota, 225 19th Avenue South, Minneapolis, MN 55455

Jonathan P. Schroeder: jps@umn.edu

Abstract

The most straightforward approaches to temporal mapping cannot effectively illustrate all potentially significant aspects of spatio-temporal patterns across many regions and times. This paper introduces an alternative approach, bicomponent trend mapping, which employs a combination of principal component analysis and bivariate choropleth mapping to illustrate two distinct dimensions of long-term trend variations. The approach also employs a bicomponent trend matrix, a graphic that illustrates an array of typical trend types corresponding to different combinations of scores on two principal components. This matrix is useful not only as a legend for bicomponent trend maps but also as a general means of visualizing principal components. To demonstrate and assess the new approach, the paper focuses on the task of illustrating population trends from 1950 to 2000 in census tracts throughout major U.S. urban cores. In a single static display, bicomponent trend mapping is not able to depict as wide a variety of trend properties as some other multivariate mapping approaches, but it can make relationships among trend classes easier to interpret, and it offers some unique flexibility in classification that could be particularly useful in an interactive data exploration environment.

Keywords

Spatio-temporal visualization; temporal mapping; census mapping; bivariate mapping

Introduction

The U.S. National Historical Geographic Information System (NHGIS) now provides online access to U.S. census summary statistics from 1790 through 2000 (<http://www.nhgis.org>, Minnesota Population Center 2004). To facilitate computer mapping and spatial analysis of these data, the NHGIS also provides digital boundary files for U.S. states, counties and census tracts as they were defined in each decennial census year. Meanwhile, several similar projects for other countries have been opening up possibilities for historical census mapping throughout the world (Knowles 2005). These new resources enable spatio-temporal analyses of population characteristics and other census-measured features at finer resolutions and over longer time periods than was previously feasible without massive investments of time and effort. Nevertheless, to investigate long-term spatio-temporal patterns in census data, analysts must still confront the challenges involved in identifying and characterizing the many complex, multidimensional patterns that can arise in a large geographic time series.

A central complicating factor is that the most straightforward approaches to visualizing geographic time series—small multiples, animated maps and change maps—cannot

effectively depict, either by themselves or in combination, all potentially significant aspects of spatio-temporal patterns. The simplest approach, chronologically-ordered small multiples of instantaneous “snapshot” maps, may require an excessively large display space or a restrictively small map scale or extent in order to present all moments in a long time series, especially when mapping multiple regions (Figure 1). Animated maps mitigate this problem by presenting a series of snapshot maps in a temporal—rather than spatial—sequence, but animation tends to be cognitively overwhelming when patterns are complex or the time series is long (Monmonier 1992, Harrower 2007), leading Dorling (1992a, 224) to suggest that, “One of the ironies of animation is that often the most effective way of seeing change over time is to not use time to show it.”

These limitations aside, small multiples and animations of snapshot maps still provide intuitive and potentially revealing illustrations of geographic time series, but any snapshot map series, whether animated or not, can depict clearly only one aspect of spatio-temporal patterns—the way spatial distributions change over time—while obscuring another important aspect—the way “temporal behaviors” vary spatially (Andrienko and Andrienko 2006). In short, this is a distinction between *trends in spatial patterns*—e.g., shifts, cluster formation, or dispersal (Figure 2a)—and *spatial patterns in trends*—e.g., pockets of similar growth rates, or rates that vary with distance from a location (Figure 2b).¹ The latter are difficult to identify in snapshot map series because doing so requires tracking changes at multiple locations across multiple map frames. Change maps, a third common approach to temporal mapping, do effectively reveal spatial patterns in trends, allowing viewers to see directly how trends vary among places but in only one regard: the magnitude of change between two moments (Figure 3). A single change map cannot depict nonlinear trend features, such as acceleration or oscillation, or the regrowth occurring in the southeastern units in Figure 2.

This paper introduces an alternative approach, bicomponent trend mapping, which depicts multiple forms of trend variation simultaneously on one map by combining the techniques of principal component analysis (PCA) and bivariate choropleth mapping. In a bicomponent trend map, two dimensions of color variation correspond to two distinct dimensions of trend variation. To facilitate the interpretation of these potentially complex color and trend variations, bicomponent trend maps employ a “bicomponent trend matrix,” which extends a standard bivariate choropleth map legend by placing a time-series graph in each color cell of the legend to indicate how colors correspond to different trend types.

There are many other multivariate mapping approaches that can also effectively illustrate spatial patterns in complex trend variations. Previous research has employed several of these, including single-component map series, with each map illustrating one trend component derived from PCA (e.g., Skaggs 1975, Eastman and Fulk 1993); time-graph maps, in which time-series graphs are used as point symbols, as in Figure 2b (e.g., Andrienko and Andrienko 2006, 2007); maps that use cluster analysis to classify trends (e.g., Hammer *et al.* 2004, Gregory 2008); and maps linked to self-organizing maps (SOMs) (e.g., Skupin and Hagelman 2005, Guo *et al.* 2006). To assess bicomponent trend mapping, I illustrate and discuss its advantages through comparisons with these other approaches.

In brief, two primary limitations of bicomponent mapping are that it depicts only two dimensions of trend variation, and it may divide natural clusters of similar trends into

¹To condense the language of Andrienko and Andrienko (2006), I use the term “trend” to indicate the way an attribute varies over time—either the complete record of the attribute at each of a series of times (the “temporal behavior”) or a synopsis of this behavior (the “temporal pattern”). The term “trend” can also be used to refer to a spatial pattern. Here, I adhere strictly to its most common use indicating temporal variation.

different classes. These limitations are not always impairments, however, and they can often be helpful. Because of them, bicomponent maps can offer some unique, useful flexibility in trend classification and can simplify map interpretation by enabling a direct, consistent correspondence between color symbol variations and trend variations. These advantages should be most useful in settings where the impairments of bicomponent mapping's limitations matter little, *i.e.*, if the two mapped components capture the variations of primary interest and if there are no strong clusters in the data.

The next section describes the time series of U.S. census tract populations used here to demonstrate visualization approaches. Later sections provide an overview of bicomponent trend mapping and then compare it with other multivariate approaches.

Study Data

Challenges in Constructing Census Time Series

An original motivation for this research was to demonstrate the breadth and potential utility of NHGIS U.S. census tract data by mapping a large subset of the data in a manner that could provide, to the extent possible, a comprehensive overview in a single display. Cartographers and analysts approaching such a goal must, however, address a number of challenges in data assembly before it is even possible to begin exploring which mapping techniques to employ. Even though historical census data are now digitally accessible, integrating censuses across time is still complicated by the many changes in how census data were collected and tabulated over time (Martin *et al.* 2002).

In the case of U.S. census tract data, there are at least three major complications for constructing time series. First, different sets of variables are tabulated in each census, and even when the same variable is available for different censuses, its exact definition may change. This restricts which variables can be meaningfully integrated into a time series and over which time spans. Second, the spatial coverage of census tract data is inconsistent over time. The U.S. Census first tabulated tract data in 1910 for only a few major cities. Coverage then generally grew to include more cities as well as more area around cities, but it remained focused on metropolitan regions until complete national coverage was provided in 1990. Constructing tract time series therefore requires a space-time tradeoff: to extend far back in time, the spatial coverage must be limited to fewer cities and fewer suburban regions; to achieve broad spatial coverage, the series must be restricted to recent decades. The third major complication in constructing tract time series is that many tract boundaries change over time, making it impossible to link one census's tract data directly to another's.

Solutions

For the present study, I address the issue of changing variable sets by examining only the total populations of census tracts. Total population has had a highly consistent definition and can therefore be relatively easily integrated and meaningfully compared over time. It is also, of course, a closely followed variable, with population growth trends being important in many settings of social science and policy analysis.

To address the issue of inconsistent historical tract coverage, I make a compromise between spatial coverage and temporal range, ultimately limiting the spatial coverage of the data to the older cores of major urban areas in order to achieve a temporal range extending from 2000 back to 1950. Specifically, I begin with the 40 most populous urban areas in 2000—those with populations greater than 900,000 (U.S. Census Bureau 2001, 2002). Then to select 2000 census tracts corresponding to these areas (given that tract boundaries do not nest within census-defined urban areas), I consider a tract to be part of an urban area if more than 90% of the tract's land area lies within the urban area *or* if more than 10% of the tract's

land area lies within the urban area *and* the tract has a 2000 density greater than 100 persons/km² (259 persons/mi²). This base set of tracts therefore covers most parts of major urban areas and omits only those parts that lie in predominantly rural tracts.

Within the 40 major urban areas, there was a rapid increase in tract data coverage prior to 1960. Of the 2000 census tracts in these areas, only 38.2% were “well covered” in 1940 (with at least 90% of their land area covered by 1940 tracts), 56.1% were well covered in 1950, and 91.4% in 1960. By 1970, the major urban areas neared complete coverage, with 98.2% of the 2000 tracts well covered. Therefore, a tract time series running back to 1970 or 1960 need not be significantly limited by historical tract coverage. However, the 1970s represent a turning point in urban population trends, as central-city declines reached a nadir and the post-war boom in inner-ring suburbs widely came to an end (Simmons and Lang 2003, Millward and Bunting 2008). To capture trends preceding this turning point adequately (and to achieve a broader demonstration of the temporal range of NHGIS tract data), the study dataset extends back to 1950, stopping short of 1940 in order to leverage the increased tract coverage between those two times.

To construct a “core set” of tracts well covered by 1950 tract data, I select from the base set the 12,711 tracts that had an estimated 1950 density greater than 200 persons/km² (518 persons/mi²) *and* which were more than 90% covered by 1950 tracts. It is important to apply the density restriction in addition to the coverage restriction in order to achieve greater consistency in the extent and type of coverage among urban areas, given that for some urban areas, 1950 tracts extended well beyond the urban fringe, while elsewhere they covered only central cities. In fact, 6 of the 40 major urban areas had no 1950 tract data and are therefore omitted, leaving only 34 urban areas in the core set. Tracts that exceed the 1950 density threshold but were not well covered by 1950 tract data represent cases of missing data for the core set, which is indicated by a light brown symbol in the maps of core tracts here.

Finally, to address the issue of changing census tract boundaries, I estimate the 1950-through-1990 populations of 2000 census tracts using areal interpolation. Specifically, I begin with *target-density weighting*, a technique that can be used to interpolate data from one census's tracts to another's (Schroeder 2007), and extend it to produce a complete time series through a procedure I term *cascading density weighting* (CDW). Using target-density weighting, 2000 census tract population densities (the “target densities”) guide the allocation of 1990 populations among 2000 census tracts, such that the estimated 1990 density distribution among the parts of 2000 tracts intersecting a 1990 tract is directly proportional to the 2000 density distribution among the intersecting tracts. CDW re-applies this procedure iteratively through time, first using the estimated 1990 population densities to guide the allocation of 1980 populations, then using estimated 1980 densities to interpolate 1970 populations, and so on in a cascading manner. For the tracts in the core dataset that were not completely covered by historical tract data (> 90% but < 100%), I use historical county data to estimate population in the uncovered part of the tract.

It may be possible to produce more accurate estimates by using advanced statistical models (e.g., Flowerdew and Green 1989, Mugglin and Carlin 1998, Kyriakidis and Yoo 2005) and/or high-resolution ancillary data, such as land cover data from satellite imagery (e.g., Langford *et al.* 1991, Holt *et al.* 2004), road network data (e.g., Xie 1995, Reibel and Bufalino 2005), or smaller-area census data (e.g., Gregory 2002). Relative to statistical modeling approaches, CDW is easier to implement, requiring only standard GIS operations, which is valuable for exploratory analyses and broad overviews. Solutions using high-resolution ancillary data also pose problems here, as there are few readily available datasets that would relate consistently well to *historical* distributions within tracts throughout the U.S., but using CDW requires nothing more than NHGIS data.

The relative simplicity of CDW also makes it possible to anticipate the types of errors that CDW is likely to produce and interpret estimates accordingly. CDW assumes that the density distribution within a tract is proportionally stable from census to census. In effect, wherever a tract was split into smaller tracts between censuses, CDW will estimate the relative (not absolute) rate of change in all of the new “split tracts” to have been exactly the same between those two censuses, and it will estimate a uniform rate of change among the split tracts in all earlier decades as well (if there were no previous boundary changes in the same area). The resulting uniformity among trends within historical tracts may hide some potentially important local trend variations and also exaggerate some variations across historical tract boundaries. But these effects are contained within relatively small areas, and broader trend patterns spanning across many historical tracts should not be greatly distorted.

Demonstration Subset

Mapping tracts throughout 34 major urban cores at an effectively legible scale requires a large display space. Therefore, to present bicomponent trend maps of the entire core set in one compilation, I have produced a small poster, presently available at http://www.pop.umn.edu/~jonathan/core_trends. The examples in this paper focus on an easier-to-map subset of three moderately large urban areas—San Diego, Minneapolis-St. Paul and St. Louis—which had 2000 populations of 2.67 million, 2.39 million and 2.08 million respectively, ranking 15th, 16th and 17th overall. Though similar in size, these three areas have experienced significantly different patterns in population trends, exemplifying well the range of trend types that occurred throughout all major urban cores. Note that the analysis below is not completely limited to this subset, as the PCA and other multivariate analyses are drawn from and apply to the complete set of tracts from all 34 urban cores. Also, the snapshot maps in Figure 1 are not limited to the core set of tracts because snapshot map series (to their advantage) can illustrate data from inconsistent spatial units without interpolation.

Mapping Principal Trend Components

Principal Component Analysis

PCA is one of the most commonly applied techniques of multivariate data analysis, and researchers have often used it to analyze geographic time-series data. To provide a few examples, Skaggs (1975) uses PCA to investigate spatio-temporal drought patterns in the U.S. through the 1930s. Eastman and Fulk (1993) use PCA to investigate spatio-temporal patterns in remotely sensed African vegetation data. Koenig and Bischoff (2004) use PCA to investigate hotel occupancy trends throughout Wales. What PCA provides in each case is a way to identify and quantify a few prominent forms of variation in trends across many places, enabling a reduction in the size and dimensionality of the analyzed data. I describe briefly here how PCA achieves this. Jolliffe (2002) and Jackson (2003) provide thorough explanations, and Johnston (1978) explains PCA in the context of geographical analysis.

PCA takes as input a multivariate dataset (a set of measurements of multiple variables for multiple observation units) and produces a set of *components*, each of which represents a single dimension in multi-dimensional attribute space, as defined by a vector of *loadings*, coefficients indicating the relationship between the component and each variable in the original dataset. The components are ordered such that the first component identifies the dimension of greatest variance in the dataset. The second component identifies the dimension of greatest variance not captured by (orthogonal to) the first component. The third component identifies the dimension of greatest variance not captured by the first and second components, and so on. There are as many output components as there are input variables, but much of the variance in a typical input dataset will be captured by the first few

principal components, making it possible to represent a dataset's key dimensions with just a few components rather than using all of the original variables.

PCA of Core Tract Population Trends

To apply PCA to the core set of tract population trends, I use as input data the *normalized rates* of population change for each of the five inter-census decades from 1950 to 2000, according to the equation:

$$\Delta y = 100 \times \left(\frac{y_{end} - y_{start}}{y_{end} + y_{start}} \right) \quad (1)$$

where Δy is the normalized rate of change between y_{start} and y_{end} , the populations at the start of and end of a decade. Normalized rates are symmetrical about 0, ranging from a minimum possible rate of -100% to a maximum of +100%, unlike conventional rates, which may range from -100% to positive infinity (Bracken and Martin 1995, Gregory and Ell 2006). This asymmetry in conventional rates generally results in positively skewed distributions with many high outliers and occasional cases of infinite growth—a significant problem for PCA, which may be overly influenced by outliers and, like most other data analysis techniques, cannot process infinite values.

It is also possible and more common to apply PCA to instantaneous measures (e.g., tract populations or densities) rather than to change rates. In the core tract dataset, as in most time series datasets, instantaneous measures are positively correlated across time; low values tend to remain relatively low, and high values tend to remain high. Applying PCA to such data produces a first component that has consistently positive (or consistently negative) loadings for all measurement times, indicating that the dimension of greatest variation is in “overall size”—how high or low values tended to be overall through time. Overall size is not of primary interest, however, when the object is to examine trends—how values tended to *change* through time.²

I also choose to apply PCA to *standardized* normalized rates. (Each decade's normalized rates are scaled and recentered to have a mean of 0 and standard deviation of 1.) This is equivalent to deriving principal components from the correlation matrix rather than the covariance matrix of the input variables. In the core tract data, the variability in normalized change rates steadily declines over time (Table 1), with the variance in 1950s rates being five times greater than the variance in 1990s rates. A PCA using unstandardized rates would therefore give five times more weight to the 1950s than to the 1990s. In contrast, using standardized rates treats the variability in each decade's rates as, in a sense, equally informative.

Table 1 summarizes the results of the PCA, and Figure 4 illustrates the trend component loadings. The first component, which accounts for 31.4% of the variance in the standardized normalized change rates, has a positive loading for each decade, with higher loadings for earlier decades, indicating that the dimension of greatest variation among tracts is the degree to which populations generally grew or declined, particularly in earlier decades (1950 to 1980). The loadings for the second component increase throughout the period, negative for the first two decades, positive for later decades, and especially high for the 1980s and 90s, indicating that the dimension of greatest variation after accounting for the first component corresponds to acceleration trends—the degree to which recent rates either exceeded or fell

²The high degree of correlation across time in many geographic time series is also a reason why it can be difficult to observe local trends in snapshot map series. Much of the apparent variation in symbols is dedicated to illustrating variations in “overall size” across space, making it harder to detect the relatively small, but still potentially important, variations across time.

below earlier change rates. The remaining three components each relate most strongly to a shift in rates between two consecutive decades: PC3 captures shifts between 1980s rates and 90s rates, PC4 captures shifts between the 70s and 80s, and PC5 captures shifts between the 50s and 60s.

Single-Component Maps

To visualize how geographic units vary in terms of principal components, a common approach is to produce a series of maps, each illustrating component scores for a single component (Figure 4). A *component score* describes an observation unit's position in attribute space in terms of distance along a principal component's axis, as determined by a linear transformation of input data using component loadings as coefficients. When all input variables have a mean of zero, as do standardized variables, a positive component score indicates deviation from the mean in the direction of the component loadings vector, a zero score indicates no deviation from the mean in that direction, and a negative score indicates deviation in the opposite direction. The mean trend is therefore an important reference point. For the core tracts, it consists of moderate growth in the 1950s, a small decline in the 70s, and smaller changes in other decades (Table 1). Each tract's component scores indicate the degree to which the tract's population trend deviates from this relatively stable mean.

The first row of maps in Figure 4 illustrates scores for the first principal component (PC1). Orange tracts generally grew and purple tracts generally declined, particularly in early decades. Accordingly, these maps closely resemble the change maps in Figure 3, which also illustrate patterns of overall growth and decline. In both cases, it is clear that inner cores generally declined while outer tracts generally grew, with San Diego distinguished by more high-growth areas, and St. Louis distinguished by more high-decline areas.

The remaining maps in Figure 4 each provide information about nonlinear trend variations that the change maps by themselves cannot reveal, beginning with the scores for PC2, which indicate the strength of recent trends particularly in relation to earlier trends. Here, orange tracts representing stronger recent trends prevail throughout San Diego's core and in the centers of Minneapolis and St. Paul (the latter indicated by two clusters of orange, one to the west in Minneapolis's center and the other to the east in St. Paul's center). Conversely, the purple tracts in the north and east of St. Louis's core and around Minneapolis indicate either deepening decline or a "cool-off" after stronger post-war trends. The maps of PC3 scores primarily indicate where and to what extent 1990s rates either exceeded 70s and 80s rates (orange) or fell behind (purple). The PC4 maps illustrate contrasts between 70s and 80s rates, and the PC5 maps illustrate contrasts between 50s and 60s rates. Because the components are ordered by decreasing variance, each row tends to illustrate weaker scores and patterns than prior rows.

Returning to the small multiples in Figure 1, it is possible to identify there some of the patterns apparent in first component scores (e.g., general growth in San Diego and general decline in St. Louis) and even, to some extent, in second component scores (e.g., continued decline in St. Louis's northern core), but overall, the patterns that are apparent in the component maps are relatively difficult to detect in snapshot maps. The small multiples convey *trends in spatial patterns* more effectively, such as the outward expansion of urban densities and the effects of rivers and mountains as obstacles to this expansion, but to identify *spatial patterns in trends*, the single-component maps are generally more effective.

It is less clear why the single-component map series would be preferable to a series of decadal change maps. Both strategies yield the same number of maps, and both would illustrate spatial patterns in trends, but the change maps would be much simpler. Mapping all components, however, is unnecessary and may even be unhelpful if the later components

are not of interest or account for very small portions of the overall variance. For example, it may be appropriate in this case to map only the first two components, as they capture 54.5% of the variance, and the later components are generally less interesting, emphasizing particular decadal shifts rather than longer-term characteristics. Mapping a small subset of components is also more obviously useful when analyzing longer time series, in which case mapping all change periods with small multiples would be difficult and an animation could be overlong.

Still, even when they are relatively short, single-component map series tend to be difficult to interpret. An analyst must first be able to relate component scores to forms of trend variation, which is demanding enough, but then, to appreciate the type of trend that occurred in any particular place requires an intricate compositing of the information drawn from each component map. For example, to interpret the trend in a place that is purple in the first row in Figure 4 and orange in the second row, one must first interpret the meaning of both of these symbols, then “sum” these interpretations (general decline plus stronger recent trends), and then add the outcome to the mean trend (from which component scores indicate deviation). This is not only challenging but also imprecise. Just how strong is the regrowth compared to the general decline? And how might the standardization of input variables affect the relationships between scores, loadings and trends?

Bicomponent Trend Matrices

A *bicomponent trend matrix* addresses some of the difficulties of interpreting trend component scores by presenting time-series graphs associated with different combinations of scores for two components (Figure 5). Each column in the matrix corresponds to a different range of scores on one component, and each row corresponds to a different range of scores on a second component. Each matrix cell contains a time-series graph (or graphs) illustrating a typical trend (or trends) for the observation units with scores in the ranges determined by the cell’s row and column.

Figure 5 presents two bicomponent trend matrices illustrating typical trends in core tract populations. In both matrices, the columns correspond to the three tertile ranges of PC1 scores (each containing one third of the core tracts), and the rows correspond to tertiles of PC2 scores. For example, the cell in the upper left of each matrix contains a graph of the typical trend for tracts in the lowest PC1 tertile *and* the highest PC2 tertile, indicating general decline with recent regrowth. (I use quantile breaks here because they are easy to interpret, they help to demonstrate the balanced symmetry of core trends, and there are no strong natural breaks to highlight. Other strategies could be more useful in other settings.)

The two matrices differ in how they present “typical” trends. Figure 5a presents mean normalized change rates by decade among tracts in each subclass, and Figure 5b presents mean mean-scaled populations by census year. (A “mean-scaled population” here is the ratio of a tract’s population in one year to the tract’s mean population over all years. Using mean-scaled populations ensures that variations in absolute population counts among tracts do not strongly influence the mean trend lines. For example, a population change of 100 in a tract with a mean population of 500 has an effect equivalent to a change of 1,000 in a tract with a mean of 5,000.) The change-rate graphs correspond usefully to the PCA input variables and to the loadings graphs in Figure 4. The graphs of mean-scaled populations, however, seem to be easier to interpret and more informative, conveying not only rates of change, indicated by line slope, but also high and low points over time. The trend graphs in later figures therefore present mean-scaled populations.

The matrices in Figure 5 encapsulate in a relatively simple graphic the two principal dimensions of variation in population trends among 12,711 tracts throughout 34 major urban

cores over 5 decades. The linear independence of principal components coupled with the tertile classification scheme helps to ensure that each matrix cell corresponds to a substantial portion of all core tracts. (The subclass sizes range from 7.1% of core tracts in the lavender class to 14.9% in the purple class.) Therefore, the matrices provide a sort of representative sample of trends across U.S. urban cores, demonstrating that the story of core population trends since 1950 is a diverse one, consisting of general decline, stability and growth in almost equal portions, coupled with similar proportions of regrowth and regression. This technique could be similarly informative when applied to data on climate, agriculture, health, retail sales, etc., in perhaps any setting where PCA is found to be useful. Even if there is no spatial or temporal dimension to the data or if there are hundreds of input variables, a generic “bicomponent matrix” could still present bar graphs of some subset of input variables. It could also present graphs of non-input variables to illustrate relationships between components and potential correlates.

A significant limitation of bicomponent trend matrices is that it may be difficult to determine how “representative” a matrix’s graphs are. The actual trends within each subclass may deviate substantially from the mean trend (or from whatever other type of “typical trend” is used). Also, the matrix presents only two components, which masks other potentially important forms of variation. A solution to the first problem is to illustrate more than just one typical trend. If there are not too many observation units, it may be possible to graph trend lines for *all* units in each subclass. Alternatively, the matrix could employ more subdivisions so that mean trends would correspond more closely to actual trends, e.g., by dividing component scores into sextiles, producing 36 subclasses (Figure 6a).

To visualize more than two components of variation, one strategy is to construct multiple bicomponent matrices, each pairing different components. It is also possible to illustrate more than two dimensions of variation in one matrix. Figure 6b presents four graphs for each PC1-PC2 tertile subclass: one illustrating the mean trend for tracts with positive PC3 and PC4 scores, one for positive PC3 and negative PC4 scores, and so on. The result is a “quadricomponent” trend matrix illustrating four principal forms of variation (though still omitting the contrasts between 1950s and 60s rates captured by PC5). Finally, to provide an indication of how prevalent each trend type is, each graph in the Figure 6 matrices also contains a bar representing the proportion of core tracts in that graph’s class.

Bicomponent Trend Maps

By applying distinct colors to the cells of a bicomponent trend matrix, as in Figures 5 and 6, it is then possible to apply the same colors in a choropleth map. This produces a *bicomponent trend map*, with a bicomponent trend matrix serving as a map legend (Figure 7a). The maps here use a diverging-diverging bivariate color scheme, as described by Brewer (1994), with a cyan/light-gray/red sequence representing PC1 and a violet/light-gray/yellow sequence representing PC2. Each color in the matrix represents a mixture of these two schemes. The upper left cell is therefore green because it combines cyan for low PC1 scores with yellow for high PC2 scores. It would also be possible to use a sequential color scheme (e.g., a light-to-dark sequence) for one or both components, but a diverging-diverging scheme is more suitable here because component scores diverge meaningfully as positive and negative numbers away from a relatively stable mean trend. The scheme used here also enables some intuitive associations, as cooler colors correspond to declining trends, warm colors correspond to growing trends, and green and yellow—often associated with spring and youth—represent trends of recent regrowth.

The bicomponent trend maps in Figure 7a directly convey several patterns in core trends that are relatively difficult, if not impossible, to detect from the snapshot map series (Figure 1), change maps (Figure 3) or single-component maps (Figure 4). The prevalence of purple in

outer cores indicates that most inner-ring suburbs have experienced low change rates after sharp post-war growth. The appearance of green in all three central business districts (CBDs) indicates that city-center redevelopment was common in recent decades. Green and gold regrowth was most prevalent within Minneapolis and St. Paul, while most of central St. Louis experienced continued decline, especially in its northern neighborhoods. Meanwhile, much of San Diego grew through all decades. The maps of all 34 urban cores reveal similar patterns as well as several other substantial differences throughout the U.S.

Comparison with Other Multivariate Approaches

Temporal Glyphs

Another approach to mapping complex trend variations is to use multivariate point symbols, or glyphs. In addition to time graphs, as used in Figure 2b and by Andrienko and Andrienko (2006, 2007), another temporal glyph is the tree-ring glyph, which depicts temporal variations through the size or shading of concentric rings (e.g., Borchert 1967 p. 326, Dorling 1992b p. 628, Harrower 2002 p. 142). Monmonier (1990) identifies several other possibilities, including clock faces and calendar glyphs. Compared to these options, trend maps using computational or statistical data reduction techniques, such as PCA and cluster analysis, are in many ways more demanding for both map producers and readers, requiring more extensive data preparation, careful consideration of different method parameters and a potentially lengthy execution time, only to produce results that are often difficult to interpret (Andrienko and Andrienko 2007). Temporal glyphs, however, often demand substantial display space to achieve legibility, especially for maps that encompass many places, as do the core tract maps. Also, even while the purely graphical nature of temporal glyphs simplifies some elements of map preparation and interpretation, it does not facilitate complementary numerical summaries or classifications in the way that most data reduction techniques do.

Multivariate Classification

Bicomponent trend maps employ a unique approach to multivariate classification, demanding a close comparison with more conventional approaches. Here I examine two: *k*-means clustering and self-organizing maps (SOMs). I first describe and demonstrate each in turn and then summarize relative strengths and weaknesses of the bicomponent approach.

K-Means Clustering—*K*-means clustering partitions observation units into *k* clusters (where *k* is a number chosen by the analyst) such that each observation unit is more similar (closer in attribute space) to the mean center of its assigned cluster than to the mean of any other cluster. The algorithm typically begins with a seed set of initial cluster centers and assigns each observation unit to the nearest center. It then recomputes the centers to equal cluster means and re-assigns units to the updated centers, repeating this process until there are no longer any changes in cluster membership. Among cluster analysis algorithms, partitioning methods such as *k*-means are generally less computationally intensive (Aldenderfer *et al.* 1984, Kaufman and Rousseeuw 2005), which has led to their use in prior research classifying and mapping long-term trends throughout large numbers of geographic units (Hammer *et al.* 2004, Gregory 2008).

The maps in Figure 7b employ the *k*-means algorithm of Hartigan and Wong (1979). The input data are, once again, standardized normalized rates of decadal population change from 1950 to 2000 among core tracts. To produce a classification similar in precision to the 9 bicomponent classes, I considered different solutions for 8 through 15 clusters. I use the 11-cluster solution because the sum of squared distances to the mean for the most dispersed

cluster declined substantially from the 8-cluster to 11-cluster solution (from 7,009 to 3,367), beyond which, improvements slowed considerably (2,627 for 15 clusters).

Figure 7b's key illustrates each cluster's mean mean-scaled population trend in an arrangement that corresponds to the bicomponent trend matrix's. Clusters are ordered horizontally by each cluster center's PC1 score and vertically by PC2 scores. Each cluster identifier consists of a number and possibly a letter. The number indicates the quadrant of the "bicomponent space" in which each cluster's center falls. Letters distinguish clusters in the same quadrant according to their sizes and their centers' distances from the mean among all tracts. Thus, among clusters centered in the second quadrant (mainly cases of decline followed by growth), Cluster 2a is the largest and closest to the mean, and Cluster 2d is the smallest and most exceptional.

Comparing the two mapping approaches in Figure 7, the similarities may be more striking than the differences. Several prominent patterns are apparent in both: decelerating growth in the outlying tracts of all cores, numerous steady-growth tracts in San Diego, steady decline in central St. Louis particularly to its north, and pervasive regrowth within the centers of Minneapolis and St. Paul. These similarities are not surprising given that both techniques rely on multivariate classification methods to encode salient trend variations. The two approaches do, however, classify trends in distinctly different ways (Figure 8), producing several meaningful differences in results.

First, cluster analysis, by definition, aims to identify real clusters in the data—groups of units that are significantly more similar to each other than they are to units in other groups. In contrast, the quantiles breaks used in the bicomponent maps is more likely to separate some very similar cases and group together very different ones. This condition may partly explain why the patterns in the bicomponent maps are somewhat more complex than those in the cluster maps. Consider, for example, the outer tracts of Minneapolis and St. Paul, where colors change frequently in the bicomponent maps while the cluster maps illustrate larger, more uniform "clusters." The implication, unsurprisingly, is that clusters in attribute space correspond to clusters in geographic space.

Another useful feature of trend clustering is that it can capture more than two dimensions of trend variation. For example, clusters 4b and 4c divide cases of strong post-war growth (with high PC1 and moderate-to-low PC2 scores) into cases of 1950s-centered growth (high PC5) and 60s-centered growth (low PC5). A bicomponent map cannot capture the same distinction because it requires partitioning on *three* components.

The cluster maps also highlight several exceptional outlier trends that the bicomponent maps do not. For example, the 2b, 2c, 2d and 3c trend clusters each contain uncommonly strong trends with sharp inflection points at different times. By isolating these trends, the cluster maps demonstrate how large-scale redevelopments causing rapid population changes have been common in CBDs and in other areas of specialized land use, as in the port-oriented and military-use tracts around San Diego Bay. This distinction between the two approaches is due in part to the operation of *k*-means, which often assigns small groups of widespread outliers to their own clusters. It is also due in part to the quantiles classification of these bicomponent maps, which masks outliers by grouping them with more typical cases.

Of course, the information gained by delineating outlier classes also comes at a cost. Each class distinguishing outliers results in fewer classes available to distinguish less extreme but more common variations. In the 11-cluster solution used here, the four smallest clusters contain only 4.4% of the tracts while the three largest contain 61%. Large portions of *k*-means cluster maps that appear to be relatively uniform—all grouped within one large cluster—may in fact contain many variations of interest that are masked by coarse

classification (which is another, less desirable condition explaining the relative simplicity of patterns in the cluster maps). To obtain a classification that gives less weight to outliers, it is possible to use other techniques of cluster analysis, such as *k*-medians as Hammer *et al.* (2004) do, or SOMs, as discussed below. One advantage of bicomponent classification in this regard is that its class breaks are flexible and do not require complete recomputation to adjust. After applying PCA once, an analyst could begin with a balanced quantiles classification and then experiment with setting breaks at more extreme values to highlight outliers. I discuss some further advantages of bicomponent maps after presenting SOMs.

Self-Organizing Maps—SOMs are a type of artificial neural network that allocates observation units to nodes spread throughout a lattice in a way that generally places units with similar characteristics near each other (Kohonen 2001, Skupin and Agarwal 2008). The result is a “mapping” of attribute space that shares many similarities with bicomponent matrices. Both techniques project multi-dimensional data onto a (usually) two-dimensional grid while also performing classification by allocating similar units to a shared cell or node, which in turn neighbors similar cells or nodes. In fact, given the relative importance of the first two components in differentiating multidimensional data, it is common for the two axes of a SOM to correspond to the first two principal components. Some SOM implementations even employ PCA to select a well-spaced set of initial node values, making a resemblance between the SOM and a bicomponent matrix even more likely.

Figure 9 presents two SOMs of the core tract population data along with a bicomponent trend matrix for comparison. Both SOMs were rotated to match the orientation of the bicomponent matrix, but other similarities are due entirely to the SOM learning process. The first SOM (Figure 9a) is a 3-by-3 square grid of nodes, producing a result that is strikingly similar to the tertile bicomponent trend matrix (Figure 9b). Choropleth maps of these SOM classes would therefore also be very similar to the bicomponent maps in Figure 7a. The 6-by-6 hexagonal SOM (Figure 9c) likewise resembles the sextile bicomponent trend matrix (Figure 6a) but presents more deviations than the 3-by-3 SOM because of its hexagonal structure and larger number of nodes.

Following the practice of Guo *et al.* (2005, 2006), the 6-by-6 SOM nodes are superimposed on a “U-matrix” shaded to indicate the degree of similarity between SOM nodes in attribute space, with darker grays indicating greater dissimilarity. Both SOMs also use a diverging-diverging bivariate color scheme similar to those used by Guo *et al.* One departure from Guo *et al.* here is the placement of trend graphs directly over SOM nodes rather than in a separate parallel coordinate plot. Also, the proportion of cases assigned to each node is represented by a bar rather than a circle.

The SOMs combine several of the advantages of cluster analysis and bicomponent matrices. As with cluster analysis, SOMs generally delineate natural clusters well and are also able to capture more than two dimensions of variation. SOMs are less sensitive to outliers than *k*-means clustering, but they are nevertheless capable of isolating exceptional clusters if they exist. SOMs improve on standard cluster analysis methods most notably by providing, like bicomponent matrices, an easily-visualized representation of class relationships with the potential added ability to indicate degrees of similarity through a U-matrix. Given these several advantages, as well as others not demonstrated here, SOMs should be a powerfully effective means of visualizing patterns in long-term trends in many settings. Nevertheless, bicomponent matrices and maps still offer some unique advantages.

Advantages of Bicomponent Classification—Many of bicomponent classification’s most useful features derive from the flexibility of its class breaks, as noted above, and from its partitioning along only two dimensions, which can simplify the representation and

interpretation of class relationships. The advantages of flexible class breaks are of greatest relevance for exploratory data analysis and for the development of interactive visualization tools. Interactive control of class breaks could facilitate investigating outliers, searching for simpler patterns, or emphasizing particular forms of variations. Interactive tools could also enable the rapid selection of different component pairings and different numbers of classes. With maps dynamically linked to a bicomponent matrix and to a scatterplot like the one in Figure 8, a user could possibly “drag” class breaks left or right, up or down, or also apply geographic brushing (Monmonier 1989, 1992). In contrast, with cluster analysis and SOMs, adjusting class membership generally requires a complete rerun using different input parameters or a different algorithm, with no simple mechanism for making fine adjustments. Meanwhile, the bicomponent approach even allows for a completely *unclassified* representation with component score magnitudes indicated by continuous color variations.

For a non-interactive display, the flexibility of class breaks is less important, and the ability of cluster analysis or SOMs to represent more than two dimensions of variation is more valuable as it enables a more complete representation in a single display. Still, achieving such “completeness” entails significant challenges in representation. In particular, it is difficult to develop effective, consistently scaled map symbols for more than a few dimensions of variation. The bivariate color schemes used here to symbolize trend clusters and SOM nodes are a reasonable strategy—possibly the best available—for distinguishing classes while also indicating which classes are similar, but bivariate schemes cannot consistently indicate *how* these classes relate because they suggest only two dimensions of variation where in fact there are several.

For example, among nodes in Figure 9c, shifts toward the right, toward redder hues, generally indicate stronger overall growth. Notice, however, that between the third, fourth and fifth trend graphs in the top row, the strongest differences appear to be in the timing of regrowth while the strength of overall growth is about the same. Here, redder hues have a different meaning than elsewhere, which a map reader can only determine through careful examination of individual trend graphs. For this SOM, producing universally consistent relationships between color variations and node properties is impossible because one cannot translate, directly and unambiguously, five possible dimensions of node variation into three-dimensional color space.

By instead illustrating only two dimensions of variation, bicomponent matrices and maps are able to achieve a direct translation of class relationships to color variations. Thus, in all bicomponent matrices and maps presented here, a shift toward a redder hue has exactly one meaning: stronger general growth, particularly in early decades. This consistency makes it possible to interpret a bicomponent map by learning and applying only four such general color rules, one for each side of the bicomponent matrix. Depending on how many classes there are, this could be substantially less demanding for map readers than memorizing or repeatedly rechecking the meaning of each symbol, as is necessary to interpret individual class relationships on a SOM or cluster map.

The orthogonality of bicomponent class divisions also helps to simplify class relationships. For example, the orange, pink, and purple bicomponent classes all have the same lower limit on PC1 scores—a division running exactly orthogonal to PC1 (Figure 8). They therefore form a coherent high-PC1 “general growth” group, distinguished from lower-PC1 classes in a consistent way, both in color and content. In addition, the differences between these three classes are as easy to interpret as their similarities: according to their distinct PC2 ranges, orange cases had the strongest recent rates relative to early rates, and purple cases had the weakest, with pink cases directly in between. Interpreting relationships among clusters or SOM nodes is more complicated because the divisions between them cut through multi-

dimensional space in ways that need not correspond to any easily interpreted dimension and cannot be represented as simply as are the bicomponent class divisions in a scatterplot, as in Figure 8.

There may still be many settings in which the liabilities of bicomponent classification outweigh its advantages. If there is no pair of meaningful principal components or if a lack of linear relationships makes a bicomponent projection unhelpful, then the simple class-color relationships are of little help. If there are some important clusters that bicomponent partitioning cannot delineate, then a clustering technique will generally be more suitable. These conditions do not appear to exist in the core tract population trends, however. The first two trend components, capturing a majority of input variance, correspond to meaningful forms of variation, and applying a simple tertile-tertile classification produces a meaningful set of prototypical trend types. Although the 3-by-3 SOM produces similarly prototypical classes, with the added benefit of being more sensitive to data clusters, the bicomponent classes have the advantages of simpler color and class relationships discussed above, with the added benefit in this case of a balanced, easy-to-interpret quantile classification scheme. Bicomponent mapping should therefore be a uniquely effective approach in this setting and may be in many others as well.

Conclusion

As we seek to make sense of the exceptionally broad and potentially complex spatio-temporal patterns encoded within new historical census data resources, or in any other large geographic time series, multivariate mapping techniques should play a central role, as they can depict spatial patterns in nonlinear trends more clearly and directly than other standard techniques of temporal mapping. The illustrations presented here indicate that, among multivariate temporal mapping approaches, bicomponent trend mapping offers several unique advantages and should be useful in many other settings.

Because this assessment is based on an application to only one study dataset and on comparisons with only a few other multivariate approaches, new insights might be gained by re-examining the technique in other settings, possibly with longer time series or with non-temporal data, and by comparing it to some other multivariate approaches, such as hierarchical cluster analysis, multidimensional scaling, or factor analysis. It should also be valuable to develop and assess an interactive data-exploration tool that employs bicomponent trend maps and matrices. It was possible here to provide only a few preliminary interpretations of the mapped patterns of core population trends. Further analysis of these trends, including comparisons with trends in other variables and the assessment of relationships with other neighborhood characteristics, might not only provide a more convincing demonstration of the utility of bicomponent mapping but also elucidate many basic patterns and processes in American urban geography.

Acknowledgments

Thanks to Martin Lacayo-Emery for timely aid in my employment of his SOM Analyst tools for ESRI ArcGIS (accessed at <http://code.google.com/p/somanalyst>, March 2010).

References

- Aldenderfer, MS.; Blashfield, RK. Sage University Paper series on quantitative applications in the social sciences, no 07-044. Beverly Hills, CA: Sage; 1984. Cluster analysis.
- Andrienko, N.; Andrienko, G. Exploratory analysis of spatial and temporal data: A systematic approach. Berlin, Germany: Springer; 2006.

- Andrienko, G.; Andrienko, N. Multimodal analytical visualization of spatio-temporal data. In: Cartwright, W.; Peterson, MP.; Gartner, G., editors. *Multimedia cartography*. Berlin, Germany: Springer; 2007. p. 327-46.
- Borchert JR. American metropolitan evolution. *Geographical Review*. 1967; 57(3):301-32.
- Bracken I, Martin D. Linkage of the 1981 and 1991 UK census using surface modeling concepts. *Environment and Planning A*. 1995; 27(3):379-90. [PubMed: 12346249]
- Brewer, CA. Color use guidelines for mapping and visualization. In: MacEachren, AM.; Taylor, DRF., editors. *Visualization in modern cartography*. Tarrytown, NY: Elsevier Science; 1994. p. 123-47.
- Dorling D. Stretching space and splicing time: From cartographic animation to interactive visualization. *Cartography & Geographic Information Systems*. 1992a; 19(4):215-27.
- Dorling D. Visualising people in space and time. *Environment and Planning B*. 1992b; 19:613-37.
- Eastman JR, Fulk M. Long time-series evaluation using standardized principal components analysis. *Photogrammetric Engineering and Remote Sensing*. 1993; 59(8):1307-12.
- Flowerdew, R.; Green, M. Statistical methods for inference between incompatible zonal systems. In: Goodchild, MF.; Gopal, S., editors. *The Accuracy of Spatial Databases*. London: Taylor & Francis; 1989. p. 239-47.
- Gregory IN. The accuracy of areal interpolation techniques: Standardizing 19th and 20th century census data to allow long-term comparisons. *Computers, Environment and Urban Systems*. 2002; 26:293-314.
- Gregory IN. Different places, different stories: Infant mortality decline in England and Wales, 1851-1911. *Annals of the Association of American Geographers*. 2008; 98(4):773-94.
- Gregory IN, Ell PS. Error-sensitive historical GIS: Identifying areal interpolation errors in time-series data. *International Journal of Geographical Information Science*. 2006; 20(2):135-52.
- Guo D, Gahegan M, MacEachren AM, Zhou B. Multivariate analysis and geovisualization with an integrated geographic visualization discovery approach. *Cartography and Geographic Information Science*. 2005; 32(2):113-32. [PubMed: 19960118]
- Guo D, Chen J, MacEachren AM, Liao K. A Visualization System for Space-Time and Multivariate Patterns (VIS-STAMP). *IEEE Transactions on Visualization and Computer Graphics*. 2006; 12:1461-74. [PubMed: 17073369]
- Hammer RB, Stewart SI, Winkler RL, Radeloff VC, Voss PR. Characterizing dynamic spatial and temporal residential density patterns from 1940-1990 across the North Central United States. *Landscape and Urban Planning*. 2004; 69(2-3):183-99.
- Harrower, MA. Doctoral dissertation. The Pennsylvania State University; 2002. Visual benchmarks: Representing geographic change with map animation.
- Harrower M. The cognitive limits of animated maps. *Cartographica*. 2007; 42(4):269-277.
- Hartigan JA, Wong MA. Algorithm AS 136: A *K*-means clustering algorithm. *Applied Statistics*. 1979; 28(1):100-8.
- Holt JB, Lo CP, Hodler TW. Dasymetric estimation of population density and areal interpolation of census data. *Cartography and Geographic Information Science*. 2004; 31(2):103-21.
- Jackson, EJ. A user's guide to principal components. Hoboken, NJ: John Wiley & Sons; 2003.
- Johnston, RJ. Multivariate statistical analysis in geography: A primer on the general linear model. London: Longman; 1978.
- Jolliffe, IT. Principal component analysis. 2. New York: Springer-Verlag; 2002.
- Kaufman, L.; Rousseuw, PJ. Finding groups in data: An introduction to cluster analysis. Reports on national historical GIS projects. In: Hoboken, NJ.; Wiley, John; SonsKnowles, AK., editors. *Historical Geography*. Vol. 33. 2005. p. 134-58.
- Koenig N, Bischoff EE. Tourism demand patterns in turbulent times: Analysing Welsh accommodation occupancy data for 1998-2001. *International Journal of Tourism Research*. 2004; 6:205-220.
- Kohonen, T. Self-organizing maps. Berlin, Germany: Springer; 2001.
- Kyriakidis PC, Yoo EH. Geostatistical prediction and simulation of point values from areal data. *Geographical Analysis*. 2005; 37(2):124-51.

- Langford, M.; Maguire, DJ.; Unwin, DJ. The areal interpolation problem: Estimating population using remote sensing in a GIS framework. In: Masser, I.; Blakemore, M., editors. *Handling Geographical Information: Methodology and Potential Applications*. London: Longman; 1991. p. 55-77.
- Martin D, Dorling D, Mitchell R. Linking censuses through time: Problems and solutions. *Area*. 2002; 34(1):82–91.
- Millward H, Bunting T. Patterning in urban population densities: A spatiotemporal model compared with Toronto 1971–2001. *Environment and Planning A*. 2008; 40(2):283–302.
- Minnesota Population Center. National Historical Geographic Information System: Pre-release Version 0.1. Minneapolis, MN: University of Minnesota; 2004.
- Monmonier M. Geographic brushing: Enhancing exploratory analysis of the scatterplot matrix. *Geographical Analysis*. 1989; 21(1):81–4.
- Monmonier M. Strategies for the visualization of geographic time-series data. *Cartographica*. 1990; 27:30–45.
- Monmonier M. Summary graphics for integrated visualization in dynamic cartography. *Cartography & Geographic Information Systems*. 1992; 19:23–36.
- Mugglin AS, Carlin BP. Hierarchical modeling in geographic information systems: Population interpolation over incompatible zones. *Journal of Agricultural, Biological, and Environmental Statistics*. 1998; 3:111–30.
- Reibel M, Bufalino ME. Street-weighted interpolation techniques for demographic count estimation in incompatible zone systems. *Environment and Planning A*. 2005; 37(1):127–39.
- Schroeder JP. Target-density weighting interpolation and uncertainty evaluation for temporal analysis of census data. *Geographic Analysis*. 2007; 39(3):311–35.
- Simmons, PA.; Lang, RE. The urban turnaround. In: Katz, B.; Lang, RE., editors. *Redefining urban & suburban America: Evidence from Census 2000*. Washington, DC: The Brookings Institution; 2003. p. 51-62.
- Skaggs RH. Drought in the United States, 1931–40. *Annals of the Association of American Geographers*. 65(3):391–402.
- Skupin, A.; Agarwal, P. Introduction: What is a self-organizing map?. In: Agarwal, P.; Skupin, A., editors. *Self-organising maps: Applications in geographic information science*. Chichester, England: John Wiley & Sons; 2008.
- Skupin A, Hagelman R. Visualizing demographic trajectories with self-organizing maps. *GeoInformatica*. 2005; 9(2):159–79.
- United States Census Bureau. [accessed December 2008.] Census 2000 summary file 1. 2001. [<http://www.factfinder.census.gov>]
- United States Census Bureau. [accessed December 2008.] UA Census 2000 TIGER/Line files. 2002. [http://www.esri.com/data/download/census2000_tigerline]
- Xie Y. The overlaid network algorithms for areal interpolation problem. *Computers, Environment, and Urban Systems*. 1995; 19(4):287–306.

Figure 1.

Snapshot map series illustrating census tract population densities in three U.S. urban areas from 1950 to 2000. The map scale must be small in order to present a complete series.

Figure 2.

Two aspects of a spatio-temporal distribution: a. the trend in spatial patterns, and b. the spatial pattern in trends.

Figure 3.
Change maps illustrating long-term population trends in 2000 census tracts for three urban cores.

Figure 4. Loadings graphs and single-component maps illustrating the principal components of normalized decadal population change rates among core tracts, 1950–2000.

Figure 5.

Two bicomponent trend matrices, each illustrating the first two principal components of normalized decadal population change rates among tracts of major urban cores, 1950–2000.

Figure 6.

Alternative bicomponent trend matrices illustrating the first two principal components of normalized decadal population change rates among tracts of major urban cores, 1950–2000: a. divided into sextiles, and b. representing two additional components (a “quadricomponent” trend matrix).

Figure 7.
Comparison of: a. bicomponent trend maps and b. trend cluster maps of core tract populations, 1950–2000.

Figure 8. Plot of all urban core tracts according to first two principal trend components. Point colors indicate trend clusters, as in Figure 7b. Frame colors indicate bicomponent classes, as in Figures 5, 6 and 7a.

Figure 9. SOMs of core tract population trends: a. a 3-by-3 rectangular SOM of standardized normalized decadal change rates, b. a tertile bicomponent matrix of the same data, and c. a 6-by-6 hexagonal SOM with a U-matrix shaded to indicate distances between nodes in attribute space.

Descriptive statistics and principal components of normalized decadal population change rates among tracts of major urban cores, 1950–2000.

Table 1

	Normalized population change rates				Pct. variance explained	Cumulative percentage
	1950s	1960s	1970s	1980s	1990s	
Mean	9.6	1.2	-4.7	-0.4	1.0	
Std. deviation	22.9	15.3	13.0	10.7	10.3	
Principal component loadings						
PC1	.55	.59	.49	.30	.10	31.4
PC2	-.36	-.29	.29	.57	.62	54.5
PC3	.20	.14	-.24	-.52	.78	72.4
PC4	.27	.09	-.77	.56	.05	87.9
PC5	.67	-.73	.14	-.01	-.01	100.0