

Spatial regionalization based on optimal information compression

Alec Kirkley ^{1,2,3}✉

Regionalization, spatially contiguous clustering, provides a means to reduce the effect of noise in sampled data and identify homogeneous areas for policy development among many other applications. Existing regionalization methods require user input such as the number of regions or a similarity measure between regions, which does not allow for the extraction of the natural regions defined solely by the data itself. Here we view the problem of regionalization as one of data compression and develop an efficient, parameter-free regionalization algorithm based on the minimum description length principle. We demonstrate that our method is capable of recovering planted spatial clusters in noisy synthetic data, and that it can meaningfully coarse-grain real demographic data. Using our description length formulation, we find that spatial ethnographic data in U.S. metropolitan areas has become less compressible over the period from 1980 to 2010, reflecting the rising complexity of urban segregation patterns in these metros.

¹Institute of Data Science, University of Hong Kong, Hong Kong, China. ²Department of Urban Planning and Design, University of Hong Kong, Hong Kong, China. ³Urban Systems Institute, University of Hong Kong, Hong Kong, China. ✉email: alec.w.kirkley@gmail.com

From the growth of economies¹ to the systemic segregation of human populations² to the environmental adaptation of ecological species³, many social and natural phenomena manifest themselves in space with high levels of clustering among similar agents or entities. Precisely defining the spatial boundaries of these clusters and observing their evolution can shed light on the fundamental processes driving the dynamics of these systems, aid in the reduction of noise in spatially sampled data^{4,5}, and facilitate the identification of regions for spatially targeted policy interventions⁶ among numerous other applications. Regionalization methods—techniques to perform spatially constrained clustering by aggregating spatial units—are typically the tools of choice for partitioning spatial data into areas of interest for such analysis. Consequently, regionalization methods have been adapted for applications across fields as diverse as climatology⁷, urban sociology⁸, hydrology⁹, geocology¹⁰, and political science¹¹.

Many approaches to regionalization typically require a significant amount of input from the user to adjust various parameters prior to performing the clustering. These tunable parameters can be used to constrain the size or shape of clusters, or to avoid crossing administrative or geographical boundaries^{12,13}. User preferences are also commonly incorporated into regionalization methods through the choice of a similarity or distance function between adjacent regions^{14,15}. Additionally, as is the case with any clustering method, a key factor existing regionalization methods consider is the choice of the number of regions, which is typically fixed by the user^{12,16} but is sometimes determined endogeneously based on user-defined thresholds for covariates of interest or other heuristics that depend on one's choice of dissimilarity between spatial units^{15,17}. An increased level of user control is desirable for many applications of regionalization, as researchers can ensure that the identified regions are suitable for the task at hand and do not violate any necessary constraints. For example, clusters extracted from regionalization methods may be used to define zones designated for different aspects of urban development, and it may be preferred that these zones do not cross significant geographical or infrastructural boundaries. In other applications of regionalization, however, such as identifying characteristic scales over which segregation or other socioeconomic phenomena persist^{18–21}, one may be interested in imposing as few assumptions as possible about how the data clusters into regions, and instead rely on the data itself to naturally define these clusters. The minimum description length (MDL) principle from information theory is a rigorous statistical framework within which one can perform inference tasks with minimal user input^{22,23}, and so provides a natural foundation for new data-driven regionalization methods.

The minimum description length principle has been applied to clustering categorical data²⁴, real-valued vector data²⁵, and other sets of objects²⁶ in spatial contexts. In ref. ²⁷, an algorithm for community detection in (aspatial) network data is proposed that identifies the partition minimizing the description length of an encoding of the network. This method, however, takes only topological information into account, which is relatively uninformative for planar networks of adjacent spatial regions (as is the case in regionalization). In ref. ²⁸, a regionalization algorithm is proposed that uses concepts from information theory to define homogeneous aggregations of spatial units, which can be identified using a greedy optimization procedure. This method works well for identifying boundaries of ethnoracial segregation, but requires the user to specify the desired number of regions and chooses the class of Bregman divergences to measure information rather than a purely combinatorial description length approach.

In this paper, we present a regionalization objective function for spatial networks with distributional metadata that is based

solely on fundamental combinatorial arguments and the minimum description length principle. By viewing the problem of regionalization from this perspective, our approach does not require the specification of any free parameters such as an explicit dissimilarity function between spatial units or a particular value for the number of regions we want the algorithm to return. Our method also takes into account the full distribution of the covariate of interest in each spatial unit, rather than summarizing each local distribution with a single statistic such as its mode, and accounts for both this spatial metadata and the topology of regional adjacencies. We describe a greedy optimization procedure used to obtain a partition of the network that approximately minimizes this description length, which involves iteratively merging the pair of nodes that maximally reduces the description length. We demonstrate our method in a series of experiments using both real and synthetic spatial data. In the first experiment, we illustrate how our method can effectively recover synthetically planted clusters in spatial distributional data, even in the presence of substantial noise. We move on to show that our method extracts meaningful regions and their evolution in real ethnoracial data by analyzing the New Haven-Milford metropolitan area of the U.S. as a case study, covering the decades between 1980 and 2010. Finally, in an experiment using a set of 110 large metropolitan areas across the U.S., we demonstrate that our method reveals the increasing complexity of urban segregation patterns over this same time period, and that this trend can be well explained by the increase in small-scale ethnoracial diversity within these metros rather than by changes in segregation patterns at large spatial scales.

Results

Cluster recovery in synthetic data. As a first test of our method, we explore its capability of recovering clusters in synthetic data. To do this, we create a synthetic model of spatial distributional data that has four tunable parameters: the number of clusters K , the number of covariate categories R , the level of statistical noise between the cluster-level distributions $\epsilon_{\text{between}}$, and the level of statistical noise within the clusters, ϵ_{within} . The model requires a spatial network $G = (V, E)$ representing the adjacencies among spatial units, and for this we use the census tract network for the New Haven-Milford metropolitan area, with $n(V) = 189$ census tracts (see Methods for details on data and mathematical variables). The specific choice of G does not tend to make a qualitative difference in the results, since the spatial networks induced by the adjacencies between units will in general have very restricted topologies²⁹. It is also possible to include variable unit populations $b(u)$ in this model, but for simplicity we set $b(u) = 10,000$ for all $u \in V$ so that these values correspond roughly to the values seen in the real U.S. census tract data used in later experiments. We show that this population heterogeneity has little effect on downstream results in Supplementary Note 3 and Supplementary Figs. 3 and 4.

To generate a realization of the model, we first randomly partition the units into contiguous clusters by picking K units (“seeds”) at random and constructing the Voronoi tessellation of the centroids of the spatial units of the network with respect to these seeds. This Voronoi tessellation places each unit into the cluster corresponding to the seed geographically nearest to the unit's centroid in terms of Euclidean distance, and in doing so tends to produce clusters that are spatially contiguous (we reject the proposed partition if it has any discontinuous partitions). The Voronoi tessellation produces relatively compact convex regions in the plane, but there are other reasonable alternative tessellations for generating the randomized contiguous partition. We denote this “planted” partition $\mathcal{P}_{\text{planted}}$ to distinguish it from the

partition \mathcal{P} inferred using our minimum description length algorithm.

Next, each cluster V_k is assigned a vector $\mathbf{x}(V_k)$, which tunes the covariate distributions within the units that comprise V_k . $\mathbf{x}(V_k)$ is drawn from a Dirichlet distribution with length- R concentration parameter $\alpha = \epsilon_{\text{between}}^{-1} \mathbf{1}_R$. This allows us to tune the level of differentiation between the cluster-level distributions, as well as the localization of these distributions. For low levels of between-cluster noise $\epsilon_{\text{between}}$ ($\epsilon_{\text{between}} \lesssim 1$), the distributions $\mathbf{x}(V_k)$ will all tend to distribute their probability relatively equally around the R categories, and there is little differentiation between the clusters V_k . On the other hand, for high levels of between-cluster noise $\epsilon_{\text{between}}$ ($\epsilon_{\text{between}} \gtrsim 5$), there will be high between-cluster variance in the distributions $\{\mathbf{x}(V_k)\}$, which will each tend to localize around a single category r . In general, the higher the between-cluster noise $\epsilon_{\text{between}}$ is, the easier it should be to recover the planted clusters in the synthetic data with our partitioning algorithm, since the clusters are more easily distinguished.

To tune the level of noise within each cluster V_k , we generate the distribution $\mathbf{x}(u) = \{b_r(u)\}_{r=1}^R / b(u)$ for each $u \in V_k$ using $\mathbf{x}(u) = (1 - \epsilon_{\text{within}})\mathbf{x}(V_k) + \epsilon_{\text{within}}\mathbf{x}_{\text{noise}}$, where $\mathbf{x}_{\text{noise}}$ is drawn from a Dirichlet distribution with concentration parameters equal to 1. If the level of within-cluster noise $\epsilon_{\text{within}} \approx 0$, then each $\mathbf{x}(u)$ for $u \in V_k$ will be roughly the same as $\mathbf{x}(V_k)$, and thus the unit-level distributions $\{b_r(u)\}$ for $u \in V_k$ are very similar. On the other hand, if the level of within-cluster noise $\epsilon_{\text{within}} \approx 1$, then the vectors $\mathbf{x}(u)$ will have high variability within the cluster V_k and the distributions $\{b_r(u)\}$ for $u \in V_k$ will share very little information. As opposed to the between-cluster noise, higher values of the within-cluster noise ϵ_{within} correspond to it being *harder* to recover the planted clusters in the synthetic data, since the unit-level distributions within each clusters are not as similar to each other. Illustrative examples of realizations of this synthetic data model used for the experiments in this section are shown in Supplementary Fig. 7.

To measure the performance of our algorithm for any particular draw from the model, we compute the normalized mutual information³⁰ between our inferred minimum description length partition \mathcal{P} and the planted partition $\mathcal{P}_{\text{planted}}$. The mutual information tells us how much information is shared between the two partitions, and its value is then normalized to fall in $[0, 1]$ so that 0 corresponds to completely uncorrelated partitions, and 1 corresponds to identical partitions (up to an arbitrary relabeling of the clusters). Letting $\mathcal{P} = \{V_k\}$ and $\mathcal{P}_{\text{planted}} = \{U_{k'}\}$, the mutual information $\text{MI}(\mathcal{P}, \mathcal{P}_{\text{planted}})$ is given by

$$\text{MI}(\mathcal{P}, \mathcal{P}_{\text{planted}}) = \sum_{k,k'} \frac{|V_k \cap U_{k'}|}{n(V)} \log \frac{n(V)|V_k \cap U_{k'}|}{|V_k||U_{k'}|}. \quad (1)$$

The mutual information can be normalized to fall in $[0, 1]$ by dividing by the average of the entropies of the individual partitions \mathcal{P} and $\mathcal{P}_{\text{planted}}$, giving

$$\text{NMI}(\mathcal{P}, \mathcal{P}_{\text{planted}}) = 2 \frac{\text{MI}(\mathcal{P}, \mathcal{P}_{\text{planted}})}{H(\mathcal{P}) + H(\mathcal{P}_{\text{planted}})}, \quad (2)$$

with

$$H(\mathcal{P}) = - \sum_k \frac{|V_k|}{n(V)} \log \frac{|V_k|}{n(V)} \quad (3)$$

and

$$H(\mathcal{P}_{\text{planted}}) = - \sum_{k'} \frac{|U_{k'}|}{n(V)} \log \frac{|U_{k'}|}{n(V)}. \quad (4)$$

The normalized mutual information is a standard and well-tested measure for comparing partitions of networks^{31,32}, but it

has a critical shortcoming for our particular application in that it gives very high baseline values to completely random contiguous partitions of spatial networks. The reason for this is that Eq. (2) compares the partitions \mathcal{P} and $\mathcal{P}_{\text{planted}}$ relative to the ensemble of all possible partitions of the network, contiguous or not, and the constraint of contiguity induces a high baseline level of correlation between the partitions. To correct for this, we rescale the normalized mutual information by subtracting off its maximum value at $\epsilon_{\text{within}} = 1$ over all simulations, which we denote $\text{NMI}_{\text{baseline}}$, and dividing by one minus this baseline value. The resulting measure is more appropriate for comparing spatially contiguous partitions, and is given by

$$\text{NMI}_{\text{rescaled}} = \frac{\text{NMI} - \text{NMI}_{\text{baseline}}}{1 - \text{NMI}_{\text{baseline}}}. \quad (5)$$

It is then easy to see when we reach the NMI value at which the partitions are minimally correlated, subject to the contiguity constraint, since the rescaled measure in Eq. (5) will be near 0. Our rescaling does not map the highest value of the NMI over the ϵ_{within} range in a given experiment to 1, so that we have better differentiation of performance in the low noise region. Indeed, we will see that the zero-noise values of the rescaled NMI are slightly less than 1 in most cases, since some sampled model realizations will by chance produce some adjacent clusters that are nearly indistinguishable.

In Fig. 1, we show the results of generating realizations of synthetic contiguous partitions from our model and running our regionalization algorithm on each of these realizations to try to recover the planted clusters. To summarize the distribution of results over the ensemble of planted partitions generated from the model, each data point represents the average rescaled normalized mutual information over 100 of these cluster recovery experiments, with error bars representing 2 standard errors in the mean. We can see that as the level of within-cluster noise ϵ_{within} increases, it becomes harder for us to recover the planted partition (as expected), but that we still have recovery better than the baseline value for reasonably high levels of within-cluster noise, for $\epsilon_{\text{between}} > 1$. (At $\epsilon_{\text{between}} = 1$, there is not enough differentiation in the latent cluster-level distributions $\mathbf{x}(V_k)$ for a distinguishable cluster structure except for at very low levels of within-cluster noise ϵ_{within} .) As expected, we can observe that the recovery task becomes easier as $\epsilon_{\text{between}}$ increases, since we have better differentiation in the latent cluster-level distributions $\{\mathbf{x}(V_k)\}$. We can see that the exact values of ϵ_{within} and $\epsilon_{\text{between}}$ at which significant enough noise is introduced to obscure the cluster structure of the data are different, since $\epsilon_{\text{within}} \in [0, 1]$ is a fractional weight and $\epsilon_{\text{between}} \in [0, \infty)$ is an inverse Dirichlet concentration parameter. Recovery performance also improves as R increases, as it is less likely for the modes of the distributions $\mathbf{x}(V_k)$ to overlap for larger R . The performance of our algorithm does not vary significantly with the number of planted clusters K , so results are displayed only for $K = 5$ for clearer visualization.

Overall, the results of Fig. 1 indicate that our minimum description length regionalization algorithm is able to successfully recover artificially planted clusters, even in the presence of substantial noise, with the performance varying as expected with the level of homogeneity within and between clusters. We now move on to examine its performance on real ethnographic distribution data.

Case study: Ethnographic composition of the New Haven-Milford metropolitan area. To illustrate how the clusters obtained with our regionalization algorithm capture meaningful

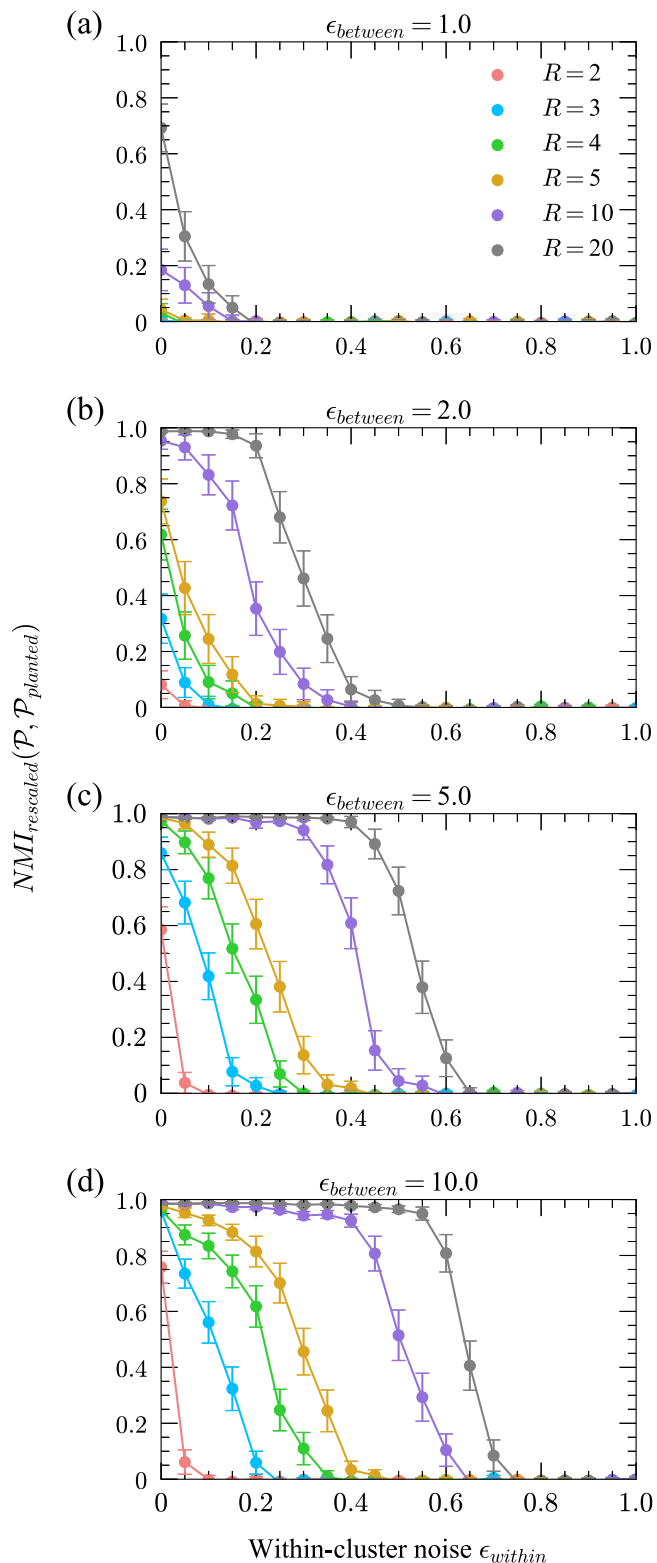


Fig. 1 Recovery of synthetic clusters. In each panel, the recovery performance of our algorithm, as measured by the rescaled normalized mutual information (NMI) of Eq. (5), is plotted on the y-axis against the level of within-cluster noise ϵ_{within} on the x-axis, for between-cluster noise **a** $\epsilon_{\text{between}} = 1.0$, **b** $\epsilon_{\text{between}} = 2.0$, **c** $\epsilon_{\text{between}} = 5.0$, and **d** $\epsilon_{\text{between}} = 10.0$. The number of covariate categories R is varied within each panel (denoted by different colors), and the number of clusters is set to $K = 5$. Error bars represent 2 standard errors in the mean.

patterns in real data, we look at a case study of the ethnoraical evolution of the New Haven-Milford, Connecticut metropolitan area, using the data described in Methods. This metro was chosen for the case study analysis due to a clearly visible spatial evolution of different ethnoraical groups and relatively low heterogeneity in census tract density in comparison with other smaller metros in our dataset, both factors allowing for a clear visual analysis of its temporal segregation patterns. Additionally, the New Haven-Milford metro exhibits a noticeable increase in ethnoraical diversity at small scales, which will help us motivate the analysis in the next section.

In Fig. 2, we show the evolution of the spatial distribution of ethnoraical groups, along with the regional boundaries inferred from minimizing the description length in Eq. (16), for the census tracts in the New Haven-Milford metro area between 1980 and 2010. Points are distributed randomly within each tract in proportion to the fraction of the population in each ethnoraical category. We can see that, in general, the clusters inferred through our algorithm correspond to heterogeneities in the spatial densities of these ethnoraical groups. The outlying tracts in the clusters, particularly in the year 2000, do not have as high a proportion of minority ethnoraical groups as the more densely packed areas of the clusters, but we can see these areas begin to fill out with minority populations over time (their inclusion status in the cluster is determined by their slightly higher relative concentrations of the minority groups dominant in the core of their cluster, compared to nearby areas).

Two emerging Black/Hispanic clusters in the north and one in the south are the primary clusters dense with minority populations that are captured by the algorithm, which assigns the rest of the metro to a single more rural/suburban and predominantly White cluster in all years (in 2000 and 2010 this cluster is broken into two due to contiguity requirements). We see that these clusters trend towards higher percentages of Hispanics relative to Non-Hispanic Blacks, which is consistent with the high influx of Latinos to the area between 1990 and 2000³³. The spatial extent of these Black/Hispanic clusters increases over time, reaching out into the less dense region of the metro that was predominantly White in 1980, which is consistent with “White flight” during deindustrialization as well as the expanding influence of Yale University in the south³⁴. In 2010, we see a slightly different configuration of clusters, with the northern Black/Hispanic clusters remaining largely intact, but the southern-most cluster splitting into a largely Black/Hispanic cluster and one relatively mixed cluster. In 2000, this mixed cluster was merged with a primarily Black cluster, but in 2010 we can see that the movement of Hispanic population into the previously Black cluster provided a high enough level of Black/Hispanic mixing to create a single dense southern-most cluster, and a separate cluster to the north with smaller overall minority populations. In 2010 we also see the emergence of a new largely Hispanic cluster to the west. These emerging clusters are reflected by an increasing optimal number of clusters, K , over the last three decades. The high level of spatial aggregation of Hispanic populations we see in the New Haven-Milford metro area is consistent with a general trend revealed by a fractal scaling analysis of large U.S. cities³⁵, which found that from 1990 to 2010 the fractal dimensions of predominantly Hispanic areas increased in most of the cities studied.

In addition to the emergent Black/Hispanic clusters at larger spatial scales, the rural/suburban tracts diversified metro-wide due to an influx of Asian and Hispanic populations to the area³⁶. For the most part these outlying tracts do not have sufficient differentiation in their ethnoraical distributions to necessitate separate clusters, and they are all grouped into a similar majority-White cluster for all four decades. However, this increasing tract-level diversity does result in

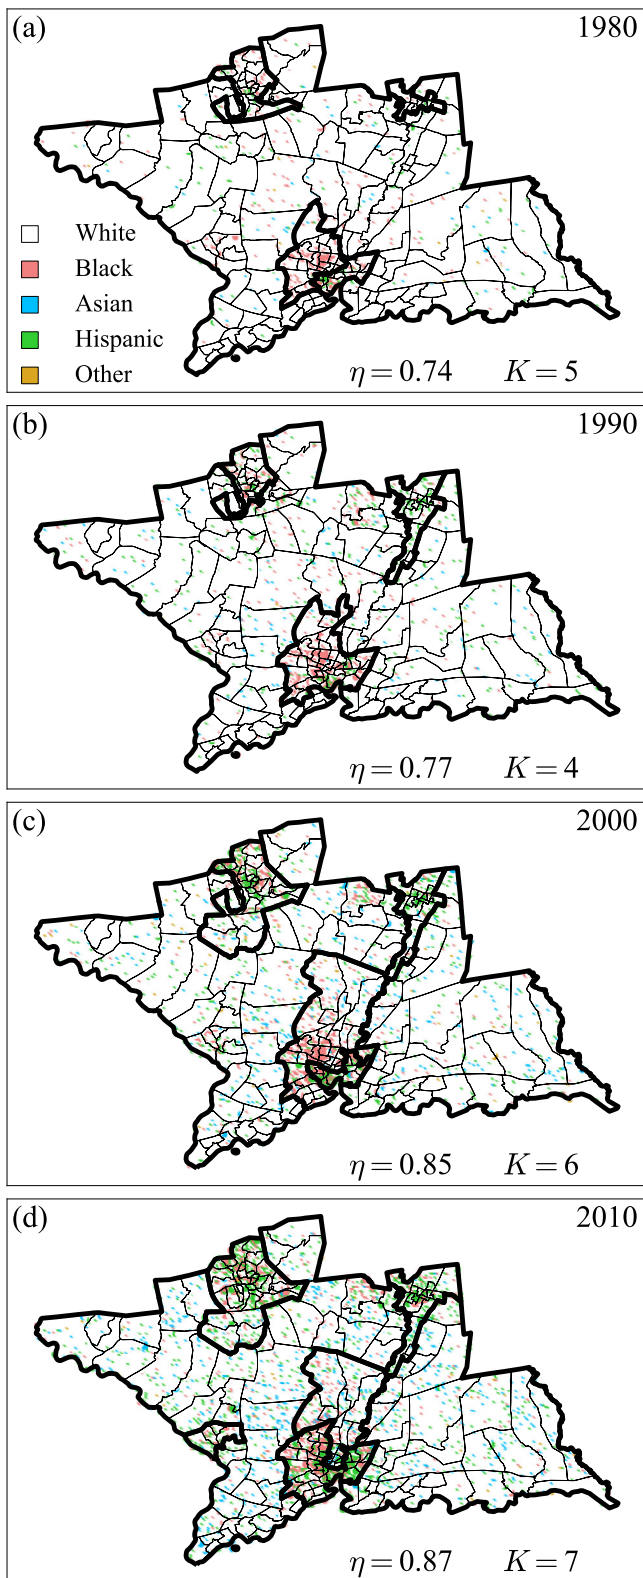


Fig. 2 Ethnoracial distributions in census tracts within the New Haven-Milford, Connecticut metropolitan area. Census tracts are delineated with thin black borders and inferred cluster boundaries from the minimum description length regionalization algorithm are shown with thick black borders for each decade **a** 1980, **b** 1990, **c** 2000, **d** 2010. Colored points are distributed at random within each tract and each color covers an area proportional to the fraction of the population within the tract that falls under the corresponding ethnoracial category. The inverse compression ratio η (Eq. (18)) and the optimal number of clusters K are shown for each decade.

greater difficulty compressing the data, as there is a clear positive trend in the inverse compression ratio η (Eq. (18)) over the four decades. We will show in the next section that a similar trend is seen across all the metros in our dataset, and that this decreasing compressibility can be better attributed to the latter effect observed in this case study (small-scale diversification) than the former (changes in large-scale segregation).

Compression of ethnoracial data across metros. Now that we have demonstrated that our regionalization method is capable of identifying meaningful clusters in ethnoracial census data, we move on to a large-scale analysis of the metro area networks described in Methods. Specifically, we look at the extent to which the data within each metro can be compressed by our algorithm according to Eq. (18), which can be used as an indicator of the overall complexity of the segregation patterns in these areas.

From a purely visual analysis, one can easily argue that the segregation patterns seen in the New Haven-Milford metro in Fig. 2 are becoming more complex over time: describing to somebody the spatial distribution of ethnoracial groups in this metro would require more effort in 2010 than in 1980. And although this concept is difficult to express in precise language due to the highly multifaceted nature of patterns in spatial data, we can capture this intuition through the inverse compression ratio of Eq. (18), which tells us how efficiently we can compress the data for its exact description to a receiver.

Despite the difficulty we may have in succinctly articulating the overall complexity of the observed segregation patterns, there are a few key features that stand out in the plots of Fig. 2. As discussed in the case study, the inverse compression ratio $\eta(D)$ increases for the New Haven-Milford over the decades spanning 1980 to 2010, and it is uncertain whether or not this increase can be better attributed to changes in tract-scale diversity or changes in large-scale segregation. The first feature of interest is the increasing diversity of a typical tract in the metro area, demonstrated by a greater and greater fraction of area covered by colored points as time progresses. The second feature that stands out is the changing spatial extent of the clustered areas, seen through the gradual absorption of the primarily White outlying tracts in 1980 into the minority-dense clusters as these clusters expand. In this section we explore the question of whether or not spatial ethnoracial patterns become more complex (as quantified by Eq. (18)) in metros other than New Haven-Milford, and to what extent the patterns we observe across these metros are consistent with each of these two features of overall diversity and changing spatial scales of clustering.

To measure the tract-level diversity of the data D in each metro area, the first feature of interest, we compute the average entropy $H_{\text{avg}}(D)$ of the ethnoracial distribution in each tract-level distribution within the metro, given by

$$\begin{aligned}
 H_{\text{avg}}(D) &= \frac{1}{n(V)} \sum_{u \in V} H(\{b_r(u)/b(u)\}) \\
 &= - \frac{1}{n(V)} \sum_{u \in V} \sum_{r=1}^R \frac{b_r(u)}{b(u)} \log \frac{b_r(u)}{b(u)},
 \end{aligned}
 \tag{6}$$

where H is the Shannon entropy. Eq. (6) will take its minimum value of 0 when the population in D is concentrated entirely into a single category r within each tract, and its maximum value of $\log R$ when all categories have equal representation in each tract within the metro.

To measure the second feature of interest, the spatial scale of clustering for a metro area, we define the characteristic cluster

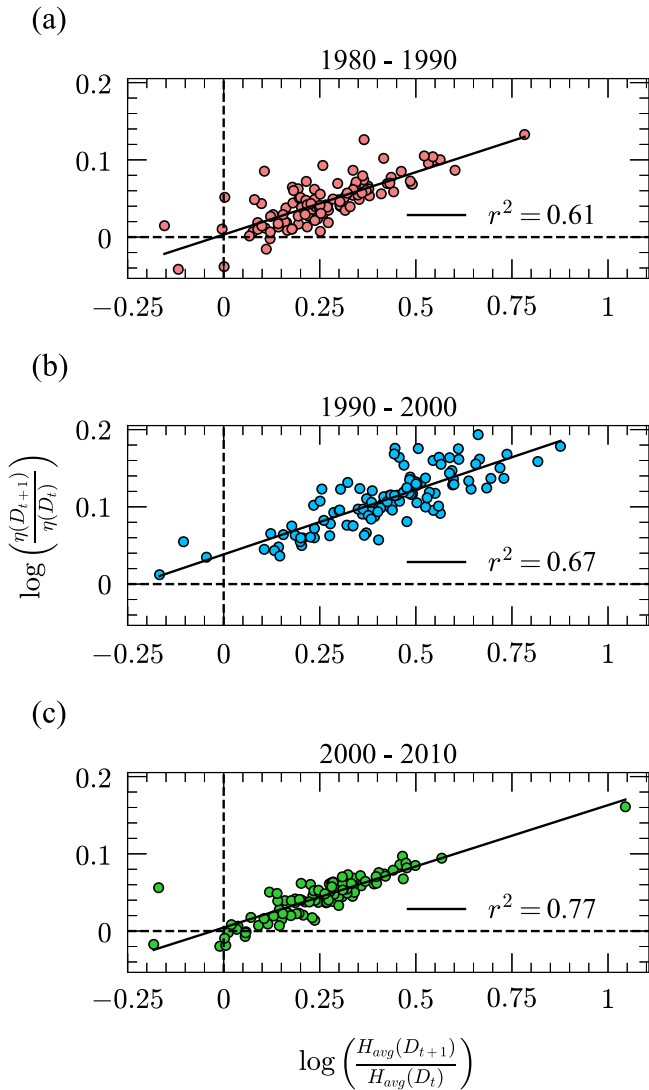


Fig. 3 Compression versus small-scale diversity. Log-ratio of consecutive inverse compression ratios η (Eq. (18)) versus the log-ratio of consecutive average tract-level diversities H_{avg} (Eq. (6)), in U.S. metros over the decades **a** $t = 1980 - 1990$, **b** $t = 1990 - 2000$, and **c** $t = 2000 - 2010$. Dotted lines at $x = 0$ and $y = 0$ are displayed for reference, along with ordinary least squares regression lines (solid black) and their coefficients of determination r^2 . The slopes of all regression lines were highly statistically significant at the 0.01 significance level. Without grouping the changes by decade, we find $r^2 = 0.70$, and that the slope is again highly significant at the 0.01 significance level.

length scale $\xi(D)$ as

$$\xi(D) = \sqrt{\frac{\sum_k A(V_k)^2}{A(V)}}, \tag{7}$$

where $A(V')$ is the area of tracts in the subset $V' \subseteq V$, and $\mathcal{P} = \{V_k\}_{k=1}^K$ is the minimum description length partition of the metro. Eq. (7) will take its minimum value of $\sqrt{A(V)/n(V)}$ when each cluster has a spatial extent of $A(V)/n(V)$ —the area of a single tract if the tracts were of equal size and each cluster only consisted of a single tract. Conversely, Eq. (7) will take its maximum value of $\sqrt{A(V)}$, the length scale of the entire metro, when the data D is best compressed with only a single cluster.

In Fig. 3, we show how changes in the inverse compression ratio η (Eq. (18)) correspond to changes in H_{avg} (Eq. (6)) across

all 110 metros for each time period in our dataset. In order to account for unobserved heterogeneity in each metro network that is constant in time—for example due to the size and topology of the metro adjacency network—as well as for potentially nonlinear dependencies, ordinary least squares (OLS) regression analysis was performed on the differences in the logarithm of each quantity over each of the periods 1980–1990, 1990–2000, and 2000–2010 (panels (a), (b), and (c), respectively, in the figure). All significance results reported in the captions hold up under Bonferroni correction for multiple comparisons³⁷.

We can see that the inverse compression ratio η is in general increasing over all time periods, as the majority of the points in Fig. 3 fall above the line $y = 0$. The average values of η over the four decades are $\{0.74, 0.77, 0.82, 0.85\}$ for $\{1980, 1990, 2000, 2010\}$. In particular, the values of η increased substantially between $t = 1990$ and $t = 2000$, with all metros in our dataset having a positive change in this quantity during this decade. This general pattern of decreasing compressibility, with the greatest change occurring during the 1990–2000 period, is consistent with the case study analysis in the previous section.

Looking at Fig. 3, we can also observe a consistently increasing level of tract-level diversity in the metro areas, as illustrated by the majority of points falling to the right of the line $x = 0$ in the three plots. The average values of H_{avg} over the four decades are $\{0.57, 0.67, 0.87, 1.02\}$ for $\{1980, 1990, 2000, 2010\}$. This observation is consistent with findings that suburbs have generally become more racially diverse³⁸, that there are an increasing number of “no-majority” communities in which no ethnoracial group makes up more than half of the population³⁹, and that the diversification of cities in the U.S. is manifested nationwide with no significant regional dependence⁴⁰. The Scranton Wilkes-Barre metro area (the rightmost point in Fig. 3) represents a clear outlier regarding changes in overall diversity, as its value of H_{avg} shot up in 2010, with roughly a 105% increase from relatively low values in the first three decades. The coefficients of determination r^2 for the regression analyses reveal that the temporal changes in H_{avg} are highly correlated with the changes in η over the same time periods, with the strongest correlation occurring between 2000 and 2010. These r^2 values, along with the statistically significant p -values of the corresponding regression line slopes (all of which had $p \ll 0.01$), suggest that the small-scale diversity within metros is an important factor for determining the complexity in segregation patterns we see according to Eq. (18).

Indeed, the results in Fig. 3 should not be too surprising: Eq. (6) has its origins in the theory of information transmission and can itself be used as a measure of spatial segregation¹⁸, like the compressibility in Eq. (18). However, Eq. (6) accounts only for diversity at small spatial scales, while the compressibility in Eq. (18) accounts for both small-scale diversity as well as large-scale homogeneity within clusters. In this way, both large-scale segregation and small-scale diversity will affect the compressibility, and therefore we need to examine both factors to determine which is a more dominant force associated with the increasing complexity we see in metros according to Eq. (18).

As shown in Supplementary Fig. 5, however, we observe no clear trend in the changes in the characteristic cluster length scales ξ (Eq (7)) across metros for each time period, with roughly half of the metros in each time period having decreasing values ξ , and half having increasing values of ξ . The metros that comprise these two halves also differ across time periods: only 18 of the 110 metros studied had monotonically increasing or decreasing values of ξ across all time periods (compared to 105 of 110 metros having a value of H_{avg} that increased throughout all decades). The r^2 values for the regression analyses in Supplementary Fig. 5 indicate that the temporal changes in ξ are poorly correlated with the changes in η over the same time periods, with r^2 values in two

of the decades even rounding to 0 up to two decimal places. The p-values corresponding to the slopes of the regression lines plotted do not indicate any statistically significant linear relationship between the plotted variables— $p = 0.52, 0.13,$ and 0.89 for panels (a), (b), and (c) respectively. These results suggest that the increasing complexity of segregation patterns we observe across metros is not substantially affected by the characteristic spatial scale at which the units can be optimally clustered in each metro (at least when considering census tracts as the fundamental unit, which obscures segregation patterns at smaller scales^{41,42}).

An important additional consideration to take into account is the effect of population, as the population in most of the metros is increasing over time, and it is reasonable to expect that this may affect the compressibility of the data. In Supplementary Fig. 6 we plot, in the same style as Fig. 3, the changes in population and the changes in compressibility of the metros over time. We can see from the OLS r^2 values that there is very little to no dependence between the population changes and the changes in compressibility of the metros (consistent with the discussion in Methods). We can also consider the effects of population and average diversity simultaneously through the following regression with city-level fixed effects

$$\log \eta_{ct} = \beta_1 \log H_{\text{avg,ct}} + \beta_2 \log b_{ct} + \alpha_c + \epsilon_{ct}, \quad (8)$$

where c and t index metros and decades respectively, $\beta_{1,2}$ are regression coefficients, α_c is an unobserved time-invariant source of heterogeneity specific to metro c (for example, based on metro c 's adjacency network topology), and ϵ is a noise term. We can then run a regression for the first differences estimators to remove the heterogeneity α_c and identify the effect $\log H_{\text{avg}}$ and $\log b$ have on $\log \eta$ when considered together. By partitioning the individual contributions of each term to the variance in the dependent variable⁴³, we find a relative importance of 97.8% for $\log H_{\text{avg}}$ versus only 2.1% for $\log b$, indicating that the average local diversity is a much more important factor for determining the compressibility than population.

Altogether, this analysis indicates that segregation patterns in large U.S. metros are becoming more complex over time from the perspective of information compression. The small-scale diversification of these metros plays an important role in increasing the complexity of these segregation patterns, while changes in population and large-scale spatial clustering among ethn racial groups are likely not major contributors.

Conclusions

Here we have presented a network regionalization algorithm based on the minimum description length principle for partitioning a set of spatial units with distributional metadata into contiguous clusters. Our method requires no user input, learning the natural clusters that result in a maximally compressed representation of the data. We demonstrate that our approach can effectively recover synthetically planted clusters in noisy spatial data and that it returns a partitioning of ethn racial census data in U.S. metropolitan areas that can allow for insights about the ethn racial segregation patterns in these metros. We find that the segregation patterns in these metros have become increasingly complex over time, in part due to the increasing small-scale ethn racial diversity of the metros over the time period studied.

There are a number of ways our method can be extended in future work. Our current formulation requires the spatial data of interest to take the form of a single discrete set of counts within each unit, but it may be possible to perform a similar description length calculation for the transmission of multiple spatial covariates simultaneously by employing the combinatorial form of the shared information between these covariates and transmitting a contingency table

indexed by groups of covariates rather than a single covariate (similar in spirit to the encoding in ref. 44). One could also develop objectives for clustering with ordinal or continuous metadata by considering the transmission on a per-symbol basis and using continuous approximations for the entropy and mutual information. This would allow us to perform regionalization with respect to a variety of attributes of interest with variable data types, for example race and income, all at once. Extension of our transmission procedure to a multi-step, hierarchical encoding scheme may also prove useful, as this would allow for multiscale regionalization. It is also possible to include additional penalties in the regionalization objective function we use in the form of Lagrange multipliers that enforce constraints on the size, shape, or populations of the clusters, which may make our method more suitable for policy-driven applications of regionalization. Additionally, using description length-based data imputation⁴⁵ one may be able to adapt our method to be robust for use with incomplete data. Finally, a comprehensive numerical comparison between the method of this paper and existing regionalization methods would shed light on the advantages and disadvantages of the MDL approach to regionalization (see Supplementary Note 4 for a qualitative comparison with similar existing methods).

Methods

Description length formulation. We represent our spatial data to be regionalized as a network $G = (V, E)$ consisting of a set of spatial units (nodes) V and a set of edges E that connect adjacent units. More precisely, the edge $(u, v) \in E$ if and only if units $u \in V$ and $v \in V$ share a length of common border. We denote the number of units in any subset $V' \subseteq V$ of the network as $n(V')$. Over this set of $n(V)$ units, there are $b(V) \geq n(V)$ individuals residing (we adopt analogous notation for $b(V')$), and each of these individuals is classified under one of R categories $r = 1, 2, \dots, R$. For example, the spatial units u that comprise the network may be census tracts or block groups, and the categories could represent race, income bracket, or occupation type. We also denote with $b_r(V')$ the number of individuals of type r in subset $V' \subseteq V$, such that $\sum_{r=1}^R b_r(V') = b(V')$.

Now, suppose we want to transmit to a receiver the entire dataset $D = \{b_r(u) : r = 1, \dots, R; u \in V\}$ consisting of the distribution of types r among individuals in all units (nodes) $u \in V$ (since we generally do not know the value r for each individual due to confidentiality concerns, these unit-level distributions are the highest granularity we consider.) We will transmit this data in multiple parts, first partitioning the units u into K disjoint, spatially contiguous clusters $\mathcal{P} = \{V_1, V_2, \dots, V_K\}$ that allow us to describe the data to the receiver at a coarse spatial scale. We then transmit the small-scale details within each of these clusters by describing how the cluster's population attributes are distributed among its individual constituent units. Our goal will be to identify a partition \mathcal{P} of the units such that most of the information we need to transmit is contained in the first part, or in other words, that the clusters describe most of the variation in the data and are internally homogeneous. Using the adjacency network representation $G = (V, E)$, we can guarantee spatial contiguity of the clusters by coarse-graining the network into super-nodes representing the clusters $\{V_k\}$ through merging nodes in V that share edges in E . A diagram of a partition \mathcal{P} of an example network and a list of the variables used in the information transmission scheme are shown in Fig. 4a.

We assume that the receiver knows there are $n(V)$ units in total that will be assigned to K clusters, and that there are $b(V)$ individuals with R distinct categories that will be assigned to units $u \in V$ (transmitting these requires a negligible amount of information, so we can safely ignore them in our description length anyway). We first need to transmit the populations $b(V_k)$ for each of the clusters V_k , which consists of a configuration of K non-negative integer values that sum to $b(V)$. Prior to transmission of the data D , we must develop a common codebook with the receiver, from which we will transmit a binary string representing the particular configuration of the populations $\{b(V_k)\}$. Assuming $K \ll b(V)$, there are approximately $(0.693)^{b(V) - 1} - 1$ possible configurations of these values we must encode, and so we will possibly have to send a bitstring of length $\lceil \log_2(0.693^{b(V) - 1} - 1) \rceil$ to the receiver to transmit the cluster-level populations $\{b(V_k)\}$. ($\lceil x \rceil$ denotes the smallest integer not less than x , and we will omit this transformation in future considerations as its contribution is negligible for $x \gg 1$). For the sake of brevity we will also denote $\log_2(x) \equiv \log(x)$. Thus, the information content (or "description length") of this step in the transmission procedure is

$$\mathcal{L}(\{b(V_k)\}) = \log \binom{b(V) - 1}{K - 1}. \quad (9)$$

Following the same logic, we can construct the description lengths for the rest of the steps required to transmit D according to this scheme. After sending the populations $\{b(V_k)\}$, we must transmit the number of units within each cluster, $\{n(V_k)\}$, for which we will construct a different codebook. This step will have a

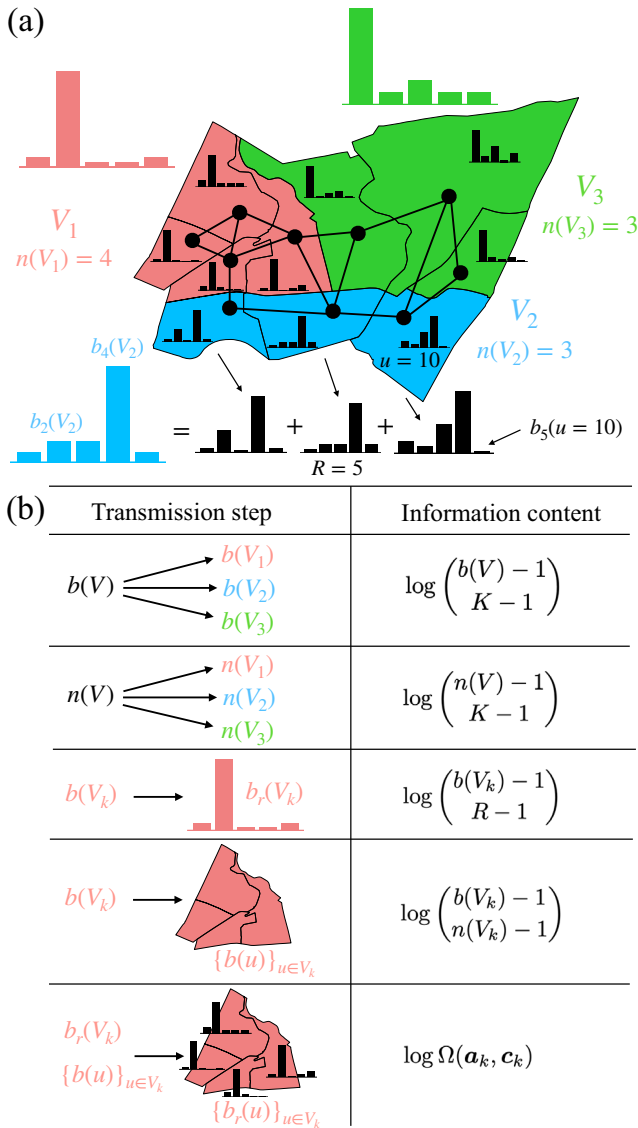


Fig. 4 Diagram of description length formulation. **a** Variables used in the description length objective (Eq. (16)), for the partition \mathcal{P} of example unit-level distributions that gives the minimum description length according to Eq. (16). The optimal contiguous partition of the underlying network of spatial adjacencies (nodes and edges in black) results in aggregated regions that capture most of the information content of the data. **b** Individual transmission steps corresponding to each of the five terms in Eq. (16), along with their corresponding information content. Arrows go from coarser objects to more detailed subsets of these objects, which requires the specification of an amount of information quantified by the term to the right of the dividing line.

description length of the same form as Eq. (9), thus

$$\mathcal{L}(\{n(V_k)\}) = \log \binom{n(V) - 1}{K - 1}. \quad (10)$$

Now, for each cluster V_k we need to transmit the size distribution $\{b_r(V_k)\}$ of categories within the population $b(V_k)$, which will have the same form as Eqs. (9) and (10). The description length of this step will be a sum over such description lengths, or

$$\mathcal{L}(\{b_r(V_k)\}) = \sum_{k=1}^K \log \binom{b(V_k) - 1}{R - 1}. \quad (11)$$

Similarly, we need to transmit the populations $b(u)$ of the units $u \in V_k$, for each cluster V_k , which will give a total description length contribution of

$$\mathcal{L}(\{b(u)\}) = \sum_{k=1}^K \log \binom{b(V_k) - 1}{n(V_k) - 1}. \quad (12)$$

The receiver now knows how many units u are in each cluster V_k , how many individuals are in each of these units, and how categories are distributed across the entire population of each V_k . The only information left to transmit is how the categories in each cluster V_k are distributed among the populations in V_k 's constituent units u . (We ignore the information required to map the final unit-level distributions to particular locations in the network.) The number of ways these values can be distributed is equivalent to the number $\Omega(\mathbf{a}_k, \mathbf{c}_k)$ of non-negative integer-valued matrices with row sums $\mathbf{a}_k = \{b(u)\}_{u \in V_k}$ and column sums $\mathbf{c}_k = \{b_r(V_k)\}_{r=1}^R$. We can see this by noting that there are $b(V_k)$ total individuals in cluster V_k , and using the identities

$$b(V_k) = \sum_{u \in V_k} b(u) \quad (13)$$

and

$$b(V_k) = \sum_{r=1}^R b_r(V_k). \quad (14)$$

The description length for this final step is thus given by

$$\mathcal{L}_{final} = \sum_{k=1}^K \log \Omega(\mathbf{a}_k, \mathbf{c}_k). \quad (15)$$

Computing $\Omega(\mathbf{a}_k, \mathbf{c}_k)$ is in general challenging, but it can be approximated in the regime $R, n(V_k) \ll b(V_k)$, which is typically the regime we encounter in practice (see ref. 44 for details on this approximation).

Taken all together, the total description length of the data D under the partition \mathcal{P} of the network G is given by the sum of Eqs (9), (10), (11), (12), and (15), thus

$$\begin{aligned} \mathcal{L}(D, \mathcal{P}) = & \log \binom{b(V) - 1}{K - 1} + \log \binom{n(V) - 1}{K - 1} \\ & + \sum_{k=1}^K \log \binom{b(V_k) - 1}{R - 1} + \sum_{k=1}^K \log \binom{b(V_k) - 1}{n(V_k) - 1} \\ & + \sum_{k=1}^K \log \Omega(\mathbf{a}_k, \mathbf{c}_k). \end{aligned} \quad (16)$$

A list of the individual transmission steps and their corresponding information content contribution to Eq. (16) is shown in Fig. 4b.

We can see that the first three terms in Eq. (16) penalize us for having a greater number of clusters K , as they will tend to contribute greater description lengths as K increases, and the fourth term will not depend on the number of clusters to first order in a Stirling approximation of the binomial coefficients. For the last term in Eq. (16), in the extreme case where there is only one category r^* that is represented in the population of the units $u \in V_k$ (i.e., $c_k[r] = 0$ for $r \neq r^*$), then we have $\Omega(\mathbf{a}_k, \mathbf{c}_k) = 1$ and the contribution from this term vanishes. More generally, there are fewer ways the categories can be distributed among the populations in V_k 's constituent units if \mathbf{c}_k is more concentrated on a single category, and so the last term in Eq. (16) will penalize us for having a high level of diversity within the clusters (or, conversely, this term encourages partitions \mathcal{P} that have homogeneous clusters).

The optimal partition $\mathcal{P} = \{V_1, \dots, V_k\}$ of the network G that minimizes the description length in Eq. (16) will allow us to communicate most of the information about the data D through the cluster-level distributions alone, but penalize us for constructing these clusters at too small a scale, since this will not save us much effort above and beyond simply transmitting all the unit-level data individually. The goal of our regionalization algorithm is to identify this partition, and we describe an algorithm to accomplish this task in the next section.

Optimization and model selection. Minimization of the description length in Eq. (16), like many other regionalization objectives¹², is a combinatorial optimization problem that can be approached in a number of ways to obtain an approximate solution. Here, we opt for a greedy solution that consists of starting with each node in its own cluster then iteratively merging the pair of adjacent clusters whose aggregation results in the largest decrease in Eq. (16), until no merges produce a negative change in the description length. We consider the two clusters of units V_k and $V_{k'}$ adjacent if and only if there exists a $u \in V_k$ and $v \in V_{k'}$ such that $(u, v) \in E$. This merging procedure thus has the benefit of naturally ensuring that the partition \mathcal{P} produces only contiguous clusters of units, since if units u and v end up in the same cluster V_k , there must be a path of edges in E that connect u and v such that all nodes along this path are also in V_k .

For any pair of clusters V_k and $V_{k'}$, we can quickly compute the change in Eq. (16) that results from their aggregation into a single cluster, $V_{k,k'}$. Supposing there are K clusters prior to the proposed merge, the change in description length from

merging V_k and $V_{k'}$ is given by

$$\begin{aligned} \Delta\mathcal{L}(k, k') = & \log\left(\frac{b(V_{k,k'}) - 1}{R - 1}\right) + \log\left(\frac{b(V_{k,k'}) - 1}{n(V_{k,k'}) - 1}\right) \\ & + \log\Omega(\mathbf{a}_{k,k'}, \mathbf{c}_{k,k'}) - \log\left(\frac{b(V_k) - 1}{R - 1}\right) \\ & - \log\left(\frac{b(V_k) - 1}{n(V_k) - 1}\right) - \log\Omega(\mathbf{a}_k, \mathbf{c}_k) \\ & - \log\left(\frac{b(V_{k'}) - 1}{R - 1}\right) - \log\left(\frac{b(V_{k'}) - 1}{n(V_{k'}) - 1}\right) \\ & - \log\Omega(\mathbf{a}_{k'}, \mathbf{c}_{k'}). \end{aligned} \quad (17)$$

Here, we have ignored the first two terms in Eq. (16), as these terms change by the same amount across all pairs k, k' and thus do not need to be computed until the optimal pair k, k' is chosen (whether or not this pair will be merged or the algorithm will terminate does depend on these first two terms, which can be computed in constant time). This expression can be evaluated in $O(n(V_k) + n(V_{k'}))$ time for each pair of clusters k, k' . Additionally, it only needs to be computed once for each pair, and can be reused for future iterations of the algorithm if the pair k, k' does not get merged (as long as each newly formed cluster gets a unique label). Once no remaining pair of clusters can be merged to reduce the description length ($\Delta\mathcal{L}(k, k') > 0$ for all adjacent pairs $V_k, V_{k'}$), the algorithm terminates.

The adjacency relations between clusters are updated as the algorithm progress by considering the clusters as “super-nodes” whose neighbor sets are merged at each step. This takes an additional $O(d_k + d_{k'})$ operations, where d_k is the number of adjacent clusters (super-nodes) to cluster (super-node) k , and is typically smaller than $O(n(V_k) + n(V_{k'}))$ for large clusters, since many clusters are adjacent to only a few others for planar graphs (this is not necessarily the case for non-planar networks). We find in practice that the algorithm scales well to large systems, running in less than order $O(n(V)^2)$ time for the entire clustering procedure (see Supplementary Note 1 and Supplementary Fig. 1).

Although the greedy algorithm used to optimize the description length in Eq. (16) has the advantages of being computationally efficient and simple to implement, it is not guaranteed to identify the true optimal partition \mathcal{P} that minimizes the description length objective over all possible partitions of the network into contiguous regions. Identifying the optimal partition \mathcal{P} is a computationally challenging optimization problem, as there are at least $O(n(V)^2)$ (and at worst exponentially many) contiguous partitions of the network one must account for⁴⁶, and even sampling such partitions is itself intractable for planar graphs⁴⁷. Additionally, fast dynamic programming approaches used for exactly solving contiguous clustering problems in one dimension are not applicable⁴⁸. However, we find in test examples that the greedy algorithm gives results quite competitive with those obtained through exhaustive enumeration of all contiguous partitions of the network to identify the true optimal partition (see Supplementary Note 2 and Supplementary Fig. 2).

The first few terms in Eq. (16) penalize us for having a large number of clusters, since we waste information describing all of the cluster-level distributions in their entirety. Meanwhile, the last term penalizes us for having a small number of clusters, since we waste information describing the small-scale details of these clusters when they encompass too broad a variety of unit-level distributions. The optimal balance, and thus the optimal value of K , lies somewhere in between with an intermediate number of clusters, and the description length in Eq. (16) thus performs model selection for K automatically. In our example applications, we therefore choose to let the description length tell us exactly how many clusters are in the data. However, in many applications it may be preferable to have a fixed value of K ¹², and this can easily be accommodated in our algorithm by simply performing the greedy merge moves until the desired number of clusters is reached.

We can assess the quality of the information compression achieved through partitioning the units into clusters by comparing the final description length $\mathcal{L}(D, \mathcal{P})$ for the optimal partition \mathcal{P} with the description length $\mathcal{L}(D, \mathcal{P}_0)$ for the trivial partition \mathcal{P}_0 in which each unit is in its own cluster (computed at the beginning of the optimization algorithm). From this we can construct an inverse “compression ratio” for the data D as

$$\eta(D) = \frac{\text{compressed size of } D}{\text{uncompressed size of } D} = \frac{\mathcal{L}(D, \mathcal{P})}{\mathcal{L}(D, \mathcal{P}_0)}. \quad (18)$$

$\eta(D)$ approaches its minimum value of 0 when the data D can be compressed extremely efficiently through partitioning the network G , and approaches its maximum value of 1 when there is no partition of G that achieves any compression of D .

Eq. (18) can thus be used as a measure of the complexity of the spatial segregation of the data D , with more complex spatial distributions of the covariate of interest resulting in higher inverse compression ratios η . Intuitively, if the data D is very easy to compress (low η), then it is highly spatially segregated into homogeneous clusters, and most of the information in D is captured at large scales. On the other hand, if the data is very hard to compress (high η), then much of the information in the data is manifested at small spatial scales, which could be due to the presence of diversity at these small spatial scales among other factors that

contribute to the multifaceted spatial nature of segregation patterns⁴⁹. The inverse compression ratio in Eq. (18) also allows us to compare the compressibility of datasets with different populations $b(V)$, numbers of categories R , number of spatial units $n(V)$, or where categories are defined differently. Indeed, for $b(V) \gg n(V) \gg R, K$ —which we typically encounter in practice for demographic data—the leading order scaling of both $\mathcal{L}(D, \mathcal{P})$ and $\mathcal{L}(D, \mathcal{P}_0)$ in Eq. (18) is $O(nR \log b)$.

Ethnoracial data in U.S. metropolitan areas. To examine the performance of our algorithm in a practical context we test our method using ethnoracial data that take the form of distributions within census tracts. Ethnoracial distributions for census tracts in U.S. metro areas were obtained from the Longitudinal Tract Database⁵⁰, which maps 2010 census tract boundaries to ethnoracial distribution data for decades going back to 1970 (data from 1970 are omitted from our analysis, as they do not include the designation of Hispanic ethnicity). The race/ethnicity categories considered are “Non-Hispanic White”, “Non-Hispanic Black”, “Asian”, “Hispanic”, and “Other”, which includes persons not categorized under the first four groups.

To process the census tract networks for each metropolitan area, we first map each census tract to its corresponding core-based Metropolitan Statistical Area (MSA) using the county designation of the tract. MSA’s are used as the metro regions for this analysis as they aim to encompass areas of unified social and economic labor market forces, while also enclosing full counties, which allows us to avoid splitting census tracts⁵¹. It is important to be mindful of this choice of metro regions, since the Modifiable Areal Unit Problem can result in different conclusions about city-level socioeconomic diversity depending on which boundaries are chosen^{52,53}.

We then use TIGER shapefile data⁵⁴ for the census tracts to determine the network $G = (V, E)$ of adjacent tracts in each MSA. Finally, the longitudinal ethnoracial distribution data is then mapped to the nodes in each network using the census tract IDs. To reduce noise as much as possible in our analysis, we kept only metros with at least 100 tracts that had complete ethnoracial distribution estimates in all tracts for the four decades 1980, 1990, 2000, and 2010. After preprocessing, 110 metro networks remained for the analysis in Results, one of which was the New Haven-Milford metro used for the case study. We make the tract adjacency networks for each metro we used in our analysis (with accompanying node metadata including ethnoracial distributions), as well as code for executing our algorithm, publicly available at https://github.com/aleckirkley/MDL_regionalization.

Data availability

All data needed to evaluate the conclusions in the paper are present in the paper or are available at https://github.com/aleckirkley/MDL_regionalization.

Code availability

The regionalization algorithm presented in this paper is available at https://github.com/aleckirkley/MDL_regionalization.

Received: 9 March 2022; Accepted: 28 September 2022;

Published online: 10 October 2022

References

1. Fujita, M., Krugman, P. R. & Venables, A. *The Spatial Economy: Cities, Regions, and International Trade* (MIT Press, 1999).
2. Brown, L. A. & Chung, S.-Y. Spatial segregation, segregation indices and the geographical perspective. *Popul. Space Place* **12**, 125–143 (2006).
3. Legendre, P. & Fortin, M. J. Spatial pattern and ecological analysis. *Vegetatio* **80**, 107–138 (1989).
4. Spielman, S. E. & Folch, D. C. Reducing uncertainty in the American Community Survey through data-driven regionalization. *PLoS ONE* **10**, e0115626 (2015).
5. Spielman, S. E. & Singleton, A. Studying neighborhoods using uncertain data from the American Community Survey: a contextual approach. *Ann. Assoc. Am. Geographers* **105**, 1003–1025 (2015).
6. Rahman, M. M. Regionalization of urbanization and spatial development: planning regions in Bangladesh. *J. Geo-Environ.* **4**, 31–46 (2004).
7. Fovell, R. & Fovell, M. Climate zones of the conterminous United States defined using cluster analysis. *J. Clim.* **6**, 2103–2135 (1993).
8. Garretton, M. & Sanchez, R. Identifying an optimal analysis level in multiscale regionalization: a study case of social distress in Greater Santiago. *Comput. Environ. Urban Syst.* **56**, 14–24 (2016).
9. Peterson, H., Nieber, J. & Kanivetsky, R. Hydrologic regionalization to assess anthropogenic changes. *J. Hydrol.* **408**, 212–225 (2011).

10. Niesterowicz, J., Stepinski, T. F. & Jasiewicz, J. Unsupervised regionalization of the United States into landscape pattern types. *Int. J. Geographical Inf. Sci.* **30**, 1450–1468 (2016).
11. George, J. A., Lamar, B. W. & Wallace, C. A. Political district determination using large-scale network optimization. *Socio-Economic Plan. Sci.* **31**, 11–28 (1997).
12. Duque, J. C., Ramos, R. & Suriñach, J. Supervised regionalization methods: a survey. *Int. Regional Sci. Rev.* **30**, 195–220 (2007).
13. Li, W., Goodchild, M. F. & Church, R. An efficient measure of compactness for two-dimensional shapes and its application in regionalization problems. *Int. J. Geographical Inf. Sci.* **27**, 1227–1250 (2013).
14. Assunção, R. M., Neves, M. C., Câmara, G. & da Costa Freitas, C. Efficient regionalization techniques for socioeconomic geographical units using minimum spanning trees. *Int. J. Geographical Inf. Sci.* **20**, 797–811 (2006).
15. Wei, R., Rey, S. & Knaap, E. Efficient regionalization for spatially explicit neighborhood delineation. *Int. J. Geographical Inf. Sci.* **35**, 135–151 (2021).
16. Aydin, O., Janikas, M. V., Assunção, R. M. & Lee, T.-H. A quantitative comparison of regionalization methods. *Int. J. Geographical Inf. Sci.* **35**, 2287–2315 (2021).
17. Duque, J. C., Anselin, L. & Rey, S. J. The max-p-regions problem. *J. Regional Sci.* **52**, 397–419 (2012).
18. Wright, R., Ellis, M., Holloway, S. R. & Wong, S. Patterns of racial diversity and segregation in the United States: 1990–2010. *Prof. Geogr.* **66**, 173–182 (2014).
19. Olteanu, M., Randon-Furling, J. & Clark, W. A. Segregation through the multiscalar lens. *Proc. Natl Acad. Sci. USA* **116**, 12250–12254 (2019).
20. Grainger, A. The role of spatial scale and spatial interactions in sustainable development. In: *Exploring Sustainable Development: Geographical Perspectives* (Earthscan, 2004).
21. Kirkley, A. Information theoretic network approach to socioeconomic correlations. *Phys. Rev. Res.* **2**, 043212 (2020).
22. Grünwald, P. D. & Grünwald, A. *The Minimum Description Length Principle* (MIT Press, 2007).
23. Cover, T. M. & Thomas, J. A. *Elements of Information Theory* (John Wiley & Sons, 2012).
24. Li, T., Ma, S. & Ogihara, M. Entropy-based criterion in categorical clustering. In: *Proceedings of the Twenty-first International Conference on Machine Learning*, p. 68, (Association for Computing Machinery, 2004).
25. Georgieva, O., Tschumitschew, K. & Klawonn, F. Cluster validity measures based on the minimum description length principle. In: *Proceedings of the International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*. 82–89 (Springer-Verlag, 2011).
26. Kirkley, A. & Newman, M. E. J. Representative community divisions of networks. *Commun. Phys.* **5**, 1–10 (2022).
27. Rosvall, M. & Bergstrom, C. T. An information-theoretic framework for resolving community structure in complex networks. *Proc. Natl Acad. Sci. USA* **104**, 7327–7331 (2007).
28. Chodrow, P. S. Structure and information in spatial segregation. *Proc. Natl Acad. Sci. USA* **114**, 11591–11596 (2017).
29. Barthélemy, M. Spatial networks. *Phys. Rep.* **499**, 1–101 (2011).
30. Vinh, N. X., Epps, J. & Bailey, J. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* **11**, 2837–2854 (2010).
31. Danon, L., Duch, J., Diaz-Guilera, A. & Arenas, A. Comparing community structure identification. *J. Stat. Mech.: Theory Exp.* **2005**, P09008 (2005).
32. Lancichinetti, A., Fortunato, S. & Radicchi, F. Benchmark graphs for testing community detection algorithms. *Phys. Rev. E* **78**, 046110 (2008).
33. Vasquez, D. W. *Latinos in New Haven, Connecticut*. (Research Report Gastón Institute Publications, No. 57, 2003).
34. Leonardo, M. D. There's no place like home: domestic domains and urban imaginaries in New Haven, Connecticut. *Identities: Glob. Stud. Cult. Power* **13**, 33–52 (2006).
35. Stepinski, T. F. & Dmowska, A. Complexity in patterns of racial segregation. *Chaos, Solitons Fractals* **140**, 110207 (2020).
36. Buchanan, M. & Abraham, M. *Understanding the Impact of Immigration in Greater New Haven* (Research Report Community Foundation for Greater New Haven, 2015).
37. Miller, R. G. *Simultaneous Statistical Inference* (Springer Verlag, 1981).
38. Orfield, M. & Luce, T. F. America's racially diverse suburbs: Opportunities and challenges. *Hous. Policy Debate* **23**, 395–430 (2013).
39. Farrell, C. R. & Lee, B. A. No-majority communities: Racial diversity and change at the local level. *Urban Aff. Rev.* **54**, 866–897 (2018).
40. Dmowska, A. & Stepinski, T. F. Spatial approach to analyzing dynamics of racial diversity in large us cities: 1990–2000–2010. *Computers, Environ. Urban Syst.* **68**, 89–96 (2018).
41. Krupka, D. J. Are big cities more segregated? Neighbourhood scale and the measurement of segregation. *Urban Stud.* **44**, 187–197 (2007).
42. Dmowska, A. & Stepinski, T. F. Improving assessment of urban racial segregation by partitioning a region into racial enclaves. *Environ. Plan. B Urban Anal. City Sci.* **49**, p. 23998083211001386 (2021).
43. Grömping, U. Relative importance for linear regression in r: the package relaimpo. *J. Stat. Softw.* **17**, 1–27 (2007).
44. Newman, M. E. J., Cantwell, G. T. & Young, J.-G. Improved mutual information measure for clustering, classification, and community detection. *Phys. Rev. E* **101**, 042304 (2020).
45. Vreeken, J. & Siebes, A. Filling in the blanks—Krimp minimisation for missing data. In: *2008 Eighth IEEE International Conference on Data Mining*. 1067–1072 (IEEE, 2008).
46. Vince, A. Counting connected sets and connected partitions of a graph. *Australas. J. Combinatorics* **67**, 281–293 (2017).
47. Najt, L., DeFord, D. & Solomon, J. Complexity and geometry of sampling connected graph partitions. Preprint <https://arxiv.org/abs/1908.08881> (2019).
48. Wang, H. & Song, M. Ckmeans.1d.dp: optimal k-means clustering in one dimension by dynamic programming. *R. J.* **3**, 29 (2011).
49. Massey, D. S. & Denton, N. A. The dimensions of residential segregation. *Soc. Forces* **67**, 281–315 (1988).
50. Logan, J. R., Xu, Z. & Stults, B. J. Interpolating US decennial census tract data from as early as 1970 to 2010: a longitudinal tract database. *Prof. Geogr.* **66**, 412–420 (2014).
51. Bettencourt, L. M. A. *Introduction to Urban Science: Evidence and Theory of Cities as Complex Systems* (MIT Press, 2021).
52. Gehlke, C. E. & Biehl, K. Certain effects of grouping upon the size of the correlation coefficient in census tract material. *J. Am. Stat. Assoc.* **29**, 169–170 (1934).
53. Cottineau, C., Hatna, E., Arcaute, E. & Batty, M. Diverse cities or the systematic paradox of urban scaling laws. *Comput. Environ. Urban Syst.* **63**, 80–94 (2017).
54. US Census Bureau. *Tiger/line Shapefiles* (US Census Bureau, 2019).

Acknowledgements

The author thanks Mark Newman for useful discussions regarding the use of mutual information for spatial clustering, and Phil Chodrow for useful discussions about the data. This work was funded by the US Department of Defense NDSEG fellowship program.

Author contributions

A.K. designed the study, performed the analyses, and wrote the manuscript.

Competing interests

The author declares no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-022-01029-4>.

Correspondence and requests for materials should be addressed to Alec Kirkley.

Peer review information *Communications Physics* thanks Christopher Lynn and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022