


STICC: a multivariate spatial clustering method for repeated geographic pattern discovery with consideration of spatial contiguity


Yuhao Kang, Kunlin Wu, Song Gao, Ignavier Ng, Jinmeng Rao, Shan Ye, Fan Zhang & Teng Fei

To cite this article: Yuhao Kang, Kunlin Wu, Song Gao, Ignavier Ng, Jinmeng Rao, Shan Ye, Fan Zhang & Teng Fei (2022) STICC: a multivariate spatial clustering method for repeated geographic pattern discovery with consideration of spatial contiguity, International Journal of Geographical Information Science, 36:8, 1518-1549, DOI: [10.1080/13658816.2022.2053980](https://doi.org/10.1080/13658816.2022.2053980)

To link to this article: <https://doi.org/10.1080/13658816.2022.2053980>

 Published online: 30 Mar 2022.

 Submit your article to this journal [↗](#)

 Article views: 697

 View related articles [↗](#)

 View Crossmark data [↗](#)



RESEARCH ARTICLE



STICC: a multivariate spatial clustering method for repeated geographic pattern discovery with consideration of spatial contiguity

Yuhao Kang^a , Kunlin Wu^b, Song Gao^a , Ignavier Ng^c, Jinmeng Rao^a,
Shan Ye^d, Fan Zhang^e  and Teng Fei^b 

^aGeoDS Lab, Department of Geography, University of Wisconsin-Madison, Madison, WI, USA; ^bSchool of Resources and Environmental Sciences, Wuhan University, Wuhan, China; ^cDepartment of Philosophy, Carnegie Mellon University, Pittsburgh, PA, USA; ^dDepartment of Geoscience, University of Wisconsin-Madison, Madison, WI, USA; ^eSenseable City Lab, Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

Spatial clustering has been widely used for spatial data mining and knowledge discovery. An ideal multivariate spatial clustering should consider both spatial contiguity and aspatial attributes. Existing spatial clustering approaches may face challenges for discovering repeated geographic patterns with spatial contiguity maintained. In this paper, we propose a Spatial Toeplitz Inverse Covariance-Based Clustering (STICC) method that considers both attributes and spatial relationships of geographic objects for multivariate spatial clustering. A subregion is created for each geographic object serving as the basic unit when performing clustering. A Markov random field is then constructed to characterize the attribute dependencies of subregions. Using a spatial consistency strategy, nearby objects are encouraged to belong to the same cluster. To test the performance of the proposed STICC algorithm, we apply it in two use cases. The comparison results with several baseline methods show that the STICC outperforms others significantly in terms of adjusted rand index and macro-F1 score. Join count statistics is also calculated and shows that the spatial contiguity is well preserved by STICC. Such a spatial clustering method may benefit various applications in the fields of geography, remote sensing, transportation, and urban planning, etc.

ARTICLE HISTORY

Received 14 August 2021
Revised 11 March 2022
Accepted 12 March 2022

KEYWORDS

Spatial clustering; spatial partition; regionalization; GeoAI; spatial contiguity

1. Introduction

A place typically has multiple features, such as environmental and socioeconomic variables. The spatial distribution of similar places may have the following two scenarios. On the one hand, as stated by the *first law of geography* that ‘near things are more related than distant things’, nearby places share similar characteristics (Tobler 1970,

Goodchild 2004, Zhu *et al.* 2018). For example, due to the nature of spatial dependence in geographic phenomena, two nearby meteorological stations may observe similar temperature, precipitation and humidity; two adjacent neighborhoods may have the same urban functions (e.g. residential area, commercial area, or educational area), because their socioeconomic characteristics are highly correlated (Yuan *et al.* 2015, Gao *et al.* 2017, Xing and Meng 2018); On the other hand, some places located in different areas may have similar attributes. For instance, Italy and California, US, have the same Mediterranean climate type; airports in two different cities are both transportation hubs. Hence, places that are not spatially adjacent to each other may still belong to the same group since their attributes are quite similar. This phenomenon, there exists places with similar attributes that are either nearby or far away, can be frequently seen on the earth.

Uncovering such patterns, which can be named as a repeated geographic pattern discovery (RGPD) problem, i.e. finding out *repeated* groups of similar places across space and maintaining the *spatial contiguity* of geographic patterns within each sub-cluster, requires multivariate spatial clustering (Murray and Estivill-Castro 1998, Miller and Han 2009). Spatial clustering aims at partitioning spatial data into a series of meaningful subgroups, and has played important roles in spatio-temporal data mining and knowledge discovery (Duque *et al.* 2007, Aldstadt 2010, Liu *et al.* 2012). By identifying spatial clusters, geographic objects with similar attributes or adjacent locations are grouped into same clusters, and are dissimilar or distant from other clusters. Detecting these clusters is necessary for a series of spatial analyses and GIS applications such as land use classification, cartographic generalization, public health, and soil mapping (Liqiang *et al.* 2013, Wang 2020, Esri 2021).

In fact, most geographic phenomena have the following two dimensions of properties, repeated patterns in attributes (single vs. repeated), and spatial contiguity (isolated vs. continuous). In this paper, *repeated patterns* refer to whether regions with similar geographic phenomena/attributes (i.e. belong to the same cluster) could appear in different locations; and the purpose of *spatial contiguity* is to assess whether and to what degree the attributes of geographic objects are spatially dependent. It should be noted that, for polygons, nearby geographic objects are their adjacent neighbors with shared boundary; while for points, Delaunay triangulation or other types of connectivity among points might be constructed first for obtaining their neighboring points. Existing spatial clustering methods can be divided into the following three categories: attribute-based clustering, regionalization-based clustering, and density-based clustering, though there are several other taxonomies (Deng *et al.* 2011). These methods are suitable for single and continuous, or repeated and isolated geographic pattern discovery. However, these methods may face challenges in handling RGPD problem, i.e. discovering repeated patterns with spatial contiguity preserved. Ideally, a multivariate spatial clustering for solving the RGPD problem should consider both spatial and aspatial attributes, that is, (1) geographic objects with similar attributes are grouped together and such groups may occur repeatedly across the space, (2) and nearby objects in physical space are encouraged to be assigned in the same cluster to maintain spatial contiguity. The two-dimensional characteristics of geographic phenomena (repeated patterns in attributes and spatial contiguity) as well as

their corresponding spatial clustering methods are plotted in Figure 1. In the following paragraphs, we describe more details for each category of spatial clustering methods. We also produce several example maps, as shown in Figure 2, based on a synthetic dataset to help demonstrate the characteristics of potential outputs by using different spatial clustering methods.

Attribute-based clustering methods, by definition, group geographic objects according to their multiple attributes. One way to perform such kind of clustering analysis is solely based on attributes while ignoring spatial relationships, such as *K-Means* (MacQueen *et al.* 1967), *BIRCH* (Zhang *et al.* 1996), *CURE* (Guha *et al.* 1998), and *SOM* (Self-organized Map) (Bação *et al.* 2005). The underlying hypothesis of these algorithms is that spatial dependence has been embodied by these multi-dimensional attributes




Figure 1. Characteristics of geographic phenomena and corresponding spatial clustering methods: horizontal-axis represents repeated patterns, and vertical-axis illustrates the degree of spatial continuity. The translucent area represent the uncertainty of locations for clustering algorithms.

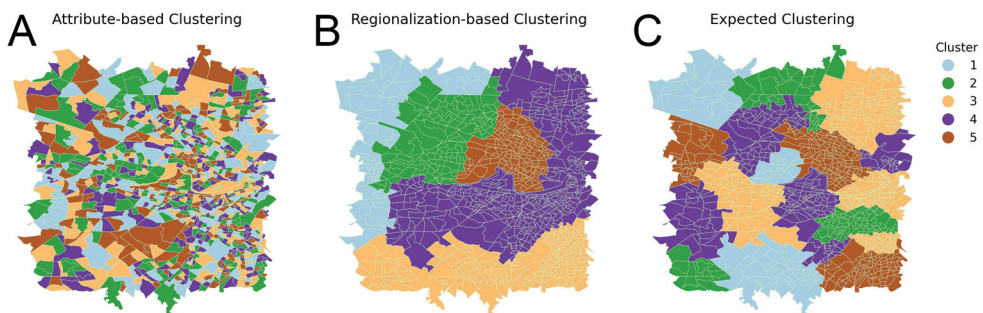


Figure 2. Example maps of spatial clustering results based on a synthetic dataset: (A) attributed-based clustering approaches; (B) regionalization-based clustering approaches; (C) expected clustering approaches (e.g. the proposed STICC) for RGPD problem. There are five clusters of regions in the study area expressed in different colors.

and thereby the spatial structures can be discovered. Another commonly used solution in practice is to treat spatial coordinates as two additional weighted attributes (Webster and Burrough 1972, Murray and Shyy 2000). It has two limitations (Perruchet 1983). First, the importance of coordinates relies on the nature of phenomenon. It is hard to decide how spatial and aspatial attributes should be combined and weighted (Duque *et al.* 2007). Second, spatial relationships, such as distant, nearby, and adjacent, which play key roles in the spatial analysis, are usually underestimated, as geography might not be a dimension in the multidimensional space (Henriques *et al.* 2009). When performing attribute-based clustering, a potential output map is shown in Figure 2(A). Geographic objects belonging to the same cluster may be distributed across space if without spatial contiguity consideration. For example, regions that belong to cluster 2 (in green) may appear in multiple locations repeatedly across the entire study area but the spatial contiguity has been destroyed. Hence, as suggested in Figure 1, attribute-based clustering may discover repeated patterns of geographic phenomenon, whereas the spatial contiguity of geographic patterns among different parts of the cluster may not be well preserved.

Density-based clustering methods such as DBSCAN (Ester *et al.* 1996), OPTICS (Ankerst *et al.* 1999), ENCLUE (Hinneburg *et al.* 1998), and ADCN (Mai *et al.* 2018), are able to find out densely located geographic patterns by examining the number of nearby geographic objects. They have been widely used in various GIS applications such as hotspot detection (Pei *et al.* 2006, Chen *et al.* 2018, Kang *et al.* 2019), urban areas of interest discovery (Hu *et al.* 2015, Liu *et al.* 2020), and taxi route and trajectory classification (Pei *et al.* 2015, Deng *et al.* 2019, Moayedi *et al.* 2019, Liu *et al.* 2021). Clusters of arbitrary shape can be discovered and the number of clusters need not to be predefined in most cases. Because density-based clustering methods usually rely on the geometric information rather than the attributes of objects (though several studies have attempted to incorporate the object attributes into these methods (Liu *et al.* 2012)), they are rarely used in finding groups of places with similar attributes. For each cluster, geographic objects gather together, and the cluster only appears once in space, while distant places with similar attributes cannot be allocated to the same cluster. Even two distant clusters with similar attributes are detected correctly, they are identified as two distinct groups. In addition, density-based clustering encounters difficulties in finding evenly distributed groups. In short, density-based clustering may not discover repeated spatial clusters, though relatively high spatial contiguity is maintained for each cluster, as shown in Figure 1. It should be noted that the clustering results using attribute-based clustering or density-based approaches may also be influenced by the degree of spatial dependence of the input geographic data. The output may still preserve the spatial contiguity to a certain degree. Such uncertainties are indicated by the translucent area in Figure 1.

The last category is regionalization-based clustering methods. Such methods are usually used for solving the p -regions problem, which is defined as the aggregation of n small areas into p geographically connected regions (Duque *et al.* 2011). It can partition space into a set of clustered regions where geographic objects are spatially connected inside those regions. Graph-based methods are commonly used for solving such a problem such as SKATER (Assunção *et al.* 2006, Aydin *et al.* 2018) and

AUTOCLUST (Aldstadt 2010). The aggregated regions determined by the regionalization-based methods are unique. Only nearby regions are aggregated to the same cluster, while distant regions with similar attributes cannot be assigned into the same cluster. As displayed in the example map in Figure 2(B), regions belonging to cluster 3 are spatially connected but only locate at the bottom of the map. The position of regionalization-based methods is thereby determined as illustrated in the Figure 1. Consequently, similar to density-based methods, spatial contiguity can be well-preserved but no repeated spatial clusters are discovered by regionalization-based methods.

In view of all of these challenges encountered in current spatial clustering methods, we aim to develop a clustering method that can consider both spatial and aspatial attributes of geographic objects. Such a method is expected to find repeated geographic patterns and maintain spatial contiguity simultaneously as shown in Figure 1. In this paper, we develop a new method that addresses these issues, which is named as *Spatial Toeplitz Inverse Covariance-based Clustering (STICC)*. The algorithm is developed to achieve the balance between multi-dimensional attributes and spatial contiguity of geographic objects with the following two characteristics. First, different from other clustering methods that usually treat each geographic object individually, we build off the algorithm using Markov random field (MRF) (Rue and Held 2005, Koller and Friedman 2009), a powerful tool that models partial correlation between different variables, to portray the dependencies of multi-dimensional attributes within a specific cluster. It builds a *subregion* between a geographic object with its nearby objects. Second, a *spatial consistency* strategy is used to encourage the nearby geographic objects to belong to the same cluster. Such a method is inspired by the *Toeplitz inverse covariance-based clustering (TICC)* method proposed by Hallac *et al.* (2017) which has been widely adopted in various multivariate time series clustering applications. It should be acknowledged that different from the time series clustering problem, the challenges of spatial clustering approaches are unique. A time-series datum only contains one-dimensional timestamp values and thereby is naturally represented as a linearly ordered sequence, while the position of each spatial object is expressed as two-dimensional coordinates and thus such a linearly ordered sequence does not naturally exist for geographic objects.

The contribution of this paper is three-fold:

1. We develop a novel spatial clustering method that considers not only the attributes of an object but also the spatial contiguity for multivariate repeated geographic pattern discovery (RGPD).
2. We validate the reliability and effectiveness of the proposed method through experiments on both synthetic examples and real-world applications.
3. We use the join count statistics to measure the spatial dependence of the clustering result for evaluating the performance of different clustering methods.

This article is structured as follows. In Section 2, we present the core idea of STICC and offer key technical implementation details, including MRF, Toeplitz matrices, and Toeplitz graphical lasso. In Section 3, we apply the method to two case studies to

validate its performance, including synthetic datasets and real-world classification scenarios. The results for different clustering methods are compared and analyzed. We make some discussions regarding hyperparameter selections of the proposed algorithm and its implications for GeoAI studies, and acknowledge limitations and opportunities for future work in [Section 4](#). Finally, we summarize and draw conclusions in [Section 5](#).

2. Method

In this section, we first outline the preliminary and required notations. We then motivate and describe the formulation of the proposed clustering problem as a mixed combinatorial and continuous optimization problem, which involves different components such as MRF, spatial Toeplitz matrix, and a penalty term to enforce spatial contiguity. Lastly, we develop an expectation-maximization (EM)-style procedure to solve the optimization problem.

It is worth noting that some part of the proposed method is inspired by the TICC method developed by Hallac *et al.* (2017). However, as described in [Section 1](#), the technical challenges are different as the TICC method is based on time series data for which a linearly ordered sequence naturally exists for the one-dimensional timestamp values, while we focus on the spatial clustering problem for which a linearly ordered sequence does not exist for the geographic objects expressed as two-dimensional coordinates.

2.1. Preliminary

The multivariate spatial clustering problem can be formulated as follows: For a given study area, clustering N geographic objects (such as points, polylines, and polygons) into K groups. The D dimensional attributes of the geographic objects are defined as

$$x = \begin{bmatrix} x_1^T \\ x_2^T \\ \vdots \\ x_N^T \end{bmatrix} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,D} \\ x_{2,1} & x_{2,2} & \dots & x_{2,D} \\ \vdots & \vdots & \vdots & \vdots \\ x_{N,1} & x_{N,2} & \dots & x_{N,D} \end{bmatrix}, \quad (1)$$

where $x_n \in \mathbb{R}^D$ is a vector that corresponds to the n th multivariate geographic object, and $x_{n,i}$ refers to the i th attribute of the n th geographic object. Furthermore, each geographic object x_n is associated with a pair of coordinates (centroids are used for polygons), denoted as $c_n = (c_{n,1}, c_{n,2})$.

Due to the existence of spatial dependence, nearby geographic objects may share similar characteristics and should be taken into account for spatial analysis. Hence, we construct a *subregion* of each point that includes several of its nearest neighbors. All objects in subregions are fed into the model for clustering. Generally, there are two ways to define a subregion, using a search radius (symmetric) or k -nearest neighbors (asymmetric) (Mai *et al.* 2018). For asymmetric definition of neighborhood, the nearest neighbors of a geographic object might not always be symmetric. For instance, a

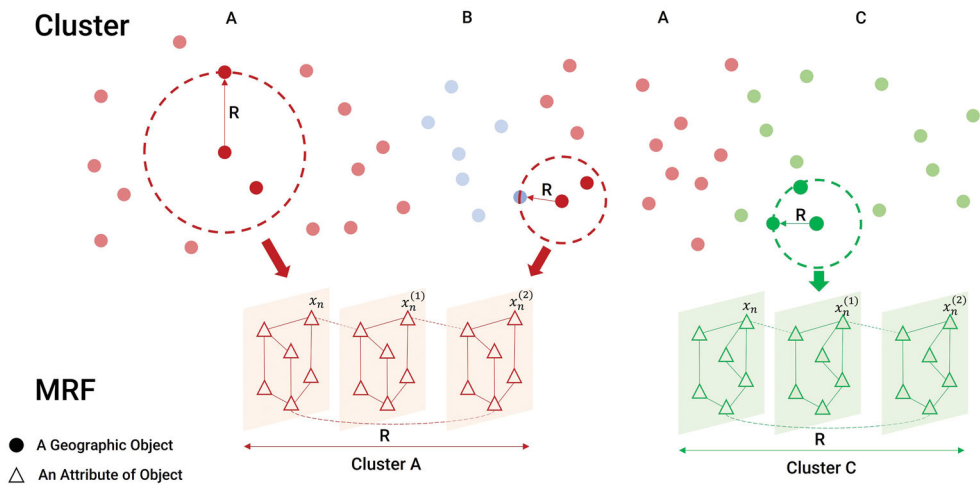


Figure 3. The idea of the proposed STICC method: geographic objects are grouped into clusters that are represented with different colors; a *subregion* is built for each geographic object with its $R-1$ nearest neighbors (here subregion radius $R=3$); x_n represents the center geographic object, and $x_n^{(r)}$ indicates its r th nearest neighbor; each cluster is characterized by an MRF to express (1) the interdependencies among attributes of different objects in the subregion (indicated by dashed lines), and (2) intradependencies among attributes of a single object (represented by solid lines); each attribute can be linked to all other attributes inside the MRF; each subregion is input into the optimization problem of STICC to learn the structure of MRF.

point A 's nearest neighbor might be B , and the nearest neighbor of B might be another point C while not A . Here, we adopt the latter one as the number of objects in each subregion is fixed, and use the center geographic object as well as its nearest $(R-1)$ neighboring objects, that is a total of R objects, for constructing the subregion. To avoid confusion, we use the term R as the 'radius' of subregions that is equivalent to the number of geographic objects in each subregion, while we use K as the number of clusters to be detected. The nearest neighbors are determined according to the spatial matrix calculated based on the pairwise distance among objects. For example, $R=1$ denotes that only the object itself is used to construct the subregion while the two nearest geographic objects of the center object are taken into account when $R=3$ as shown in Figure 3.

We then concatenate each object x_n and its $R-1$ nearest objects $x_n^{(1)}, \dots, x_n^{(R-1)}$, sorted in ascending order based on distance, into a vector; specifically, we write $X_n = (x_n, x_n^{(1)}, \dots, x_n^{(R-1)}) \in \mathbb{R}^{DR}$ as the subregion in which x_n is the centering geographic object. In addition, denote by $X_n^{(1)}$ the nearest subregion of X_n . The proposed STICC performs clustering on these stacked subregions but not on the set of all objects directly. This is different from the case of TICC that focuses on time series data, because a linearly ordered sequence naturally exists for time series and therefore the time window adopted by TICC is symmetric. The concept of time window for time series is analogous to our definition of subregions for spatial data, though the latter is not always symmetric.

2.2. Spatial Toeplitz inverse covariance-based clustering

2.2.1. Representing dependencies among attributes with Markov random field

Places with similar attributes should be grouped into the same cluster. To represent dependencies among multiple attributes of geographic objects, a Markov random field (MRF) is constructed for each cluster, in which each node in the network corresponds to an attribute (variable), and each edge indicates the dependency between different attributes (variables). Each geographic object is represented as a layer in MRF. The advantages of using MRF are two-fold.

First, edges in an MRF can be loosely interpreted as the partial correlations of a given pair of variables conditioned on the remaining variables (Rue and Held 2005, Koller and Friedman 2009). That is, any pair of variables in the network is non-adjacent if their computed partial correlation is zero, which means that an edge exists between two variables only if they are conditionally dependent given the remaining variables. Therefore, MRF is a powerful tool for modeling dependencies between different variables, as partial correlations control for the effect of the potential confounding variables. This is unlike the standard correlation method that does not take confounders into account.

Second, when using the MRF network, how a variable may affect other variables, conditioning on the remaining variables, is illustrated through its adjacencies, which provides interpretable insights to demonstrate the characteristics of clusters. For instance, assume that we want to identify urban functional zones in cities, the MRF structure of a cluster that is identified as 'commercial regions' in a city may illustrate the partial correlations between two variables, e.g. the number of shopping malls and the number of restaurants, in a subregion. Suppose in the MRF structure that we identify a high partial correlation between the number of shopping malls of the center geographic object and the number of restaurants of its nearest geographic object. This implies that after controlling all the other variables, such as number of residential buildings, number of schools, number of hospitals, etc., these two variables are highly correlated. It is worth noting that there is a distinction between this notion of partial correlation and the standard correlation. The former one, as described, takes into account the confounding variables, and therefore may be less susceptible to spurious relationship.

As mentioned above, the *subregions* of each geographic object are constructed for spatial clustering with consideration of spatial dependence, which serve as basic units for the MRFs. Hence, such MRF structures could have multiple layers (i.e. each layer represents a geographic object, and its multivariate values and their dependencies are represented as graph nodes and edges respectively) defined by the number of objects R in subregions. The edges of graphs are both within a layer and across different layers, which correspond to the intradependencies and interdependencies of the object attributes, respectively. An example is shown in Figure 3, in which the cluster's MRF contains three layers that represents the geographic object itself, the nearest object, and the second nearest object, respectively. For each cluster identified, the partial correlation structure of all geographic objects inside subregions of this cluster is depicted. It should be noted that each subregion is grouped purely based on the attribute dependency structure of the cluster to which the center geographic object

belongs. Hence, the clusters are spatial-invariant. That is to say, when assigning each geographic object to the cluster, the starting position of each subregion does not matter.

For each cluster, since the dependencies between attributes (variables) are different from other clusters, the MRF network structure is also different. For the k th cluster, we define its structure of MRF using the inverse covariance matrix of a multivariate Gaussian distribution, denoted as $\Theta_k \in \mathbb{R}^{DR \times DR}$. As described, the inverse covariance matrix illustrates the conditional independency structures among attributes inside a subregion. By definition, if the entry $(\Theta_k)_{i,j} = 0$, then these two attributes i and j are conditionally independent given the values of all other attributes in the subregion.

Next, we will introduce the specific structure of the inverse covariance matrix that helps preserve spatial invariance, including the *spatial Toeplitz matrix* in Section 2.2.2 and a *Toeplitz graphical lasso* strategy to estimate the inverse covariance matrix in Section 2.3.2.

2.2.2. Spatial Toeplitz matrix

To help preserve spatial invariance within a subregion, we restrict the inverse covariance matrices to follow the block Toeplitz form (i.e. a special diagonal-constant matrix) (Akaike 1973). In particular, the $DR \times DR$ inverse covariance matrix of each cluster is defined as

$$\Theta_k = \begin{bmatrix} A_k^{(0)} & (A_k^{(1)})^T & (A_k^{(2)})^T & \dots & \dots & (A_k^{(R-1)})^T \\ A_k^{(1)} & A_k^{(0)} & (A_k^{(1)})^T & \ddots & & \vdots \\ A_k^{(2)} & A_k^{(1)} & A_k^{(0)} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & (A_k^{(1)})^T & (A_k^{(2)})^T \\ \vdots & & \ddots & A_k^{(1)} & A_k^{(0)} & (A_k^{(1)})^T \\ A_k^{(R-1)} & \dots & \dots & A_k^{(2)} & A_k^{(1)} & A_k^{(0)} \end{bmatrix}, \quad (2)$$

where $A_k^{(0)}, A_k^{(1)}, \dots, A_k^{(R-1)} \in \mathbb{R}^{D \times D}$. Both *intradependencies* and *interdependencies* are indicated by the Toeplitz matrix. The former refers to the sub-block $A_k^{(0)}$ whose entry $(A_k^{(0)})_{i,j}$ represents the partial correlation between attributes i and j in the same geographic object within the k th cluster. Additionally, the off-diagonal sub-blocks refer to the relationships between the attributes of different geographic objects in the sub-region. For example, suppose a geographic object x_n belong to the k th cluster. Then, $(A_k^{(1)})_{i,j}$ indicates the relationship between the attribute i of x_n and the attribute j of its nearest neighbor $x_n^{(1)}$; similarly, $A_k^{(r)}, 1 < r < R$ shows the edge structure between the geographic object x_n and its r th nearest neighbor $x_n^{(r)}$. Using the block Toeplitz structure of the inverse covariance, we could characterize the relationships between multiple attributes across space with the spatial-invariance assumption that each geographic object only depends on its $(R-1)$ nearest neighbors regardless of its absolute location. Note that such attribute dependencies may vary across space but we assume it is relatively stable within each cluster.

2.2.3. Overall optimization problem

After introducing the MRF and the Toeplitz matrix, the overall spatial clustering problem can be defined as an optimization problem. The objective is to solve for these K inverse covariances $\Theta = \{\Theta_1, \dots, \Theta_K\}$ for all clusters, and the assignment sets for the geographic objects $\mathbf{P} = \{P_1, \dots, P_K\}$ with $P_i \subset \{1, 2, \dots, N\}$. It is termed as the *spatial Toeplitz inverse covariance-based clustering* (STICC), which solves the following optimization problem:

$$\min_{\Theta \in \mathcal{T}, \mathbf{P}} \sum_{k=1}^K \left[\sum_{X_n \in P_k} \left(\underbrace{-\mathcal{L}(\Theta_k; X_n)}_{\text{log likelihood}} + \beta \underbrace{\mathbb{1}\{X_n^{(1)} \notin P_k\}}_{\text{spatial consistency}} \right) + \underbrace{\|\lambda \odot \Theta_k\|_{1, \text{off}}}_{\text{sparsity}} \right], \quad (3)$$

where \mathcal{T} denotes the set of $DR \times DR$ matrices that are symmetric block Toeplitz. For each geographic object x_n , recall that X_n is defined as the subregion in which x_n is the center point. In the optimization problem, each subregion X_n is assigned to one cluster, and $-\mathcal{L}(\Theta_k; X_n)$ denotes the negative log likelihood that the subregion X_n belongs to the k th cluster. The cluster assignment of the subregion X_n is used as that of the geographic object x_n .

Here, $\mathbb{1}\{X_n^{(1)} \notin P_k\}$ is a spatial consistency indicator function that determines if X_n and its nearest subregion $X_n^{(1)}$ are in the same group, defined as

$$\mathbb{1}\{X_n^{(1)} \notin P_k\} = \begin{cases} 0, & \text{if } X_n, X_n^{(1)} \in P_k, \\ 1, & \text{otherwise.} \end{cases} \quad (4)$$

β is a hyperparameter that controls the importance of this term. This penalty term is used to encourage the neighboring subregions X_n and $X_n^{(1)}$ to be assigned to the same cluster, so that the spatial contiguity is maintained.

An ℓ_1 penalty term $\|\lambda \odot \Theta_k\|_{1, \text{off}}$ is incorporated to enforce the sparsity on the off-diagonal entries of the inverse covariance matrices, where $\lambda \in \mathbb{R}^{DR \times DR}$ is a hyperparameter that controls the sparsity level, and \odot denotes the entry-wise product. This penalty term helps enforce sparsity on the edges of the corresponding MRFs, since the non-zero entries in the inverse covariance matrices correspond to the underlying structure of the MRFs. Note that the sparsity assumption, or more specifically, the ℓ_1 constraint, has been adopted and is essential in many areas such as compressed sensing (Candes and Tao 2005, Donoho 2006), linear regression (Tibshirani 1996, Zou 2006), and graphical model selection (Friedman *et al.* 2008, Aragam and Zhou 2015, Ng *et al.* 2020). This is especially useful for high-dimensional tasks with a large number of variables and limited samples (Hastie *et al.* 2015), as is the case in our multivariate spatial clustering.

Since we focus on the Gaussian inverse covariance matrices and assume that the joint distribution of X_n follows a multivariate Gaussian distribution, the log likelihood $\mathcal{L}(\Theta_k; X_n)$ is given by

$$\mathcal{L}(\Theta_k; X_n) = -\frac{1}{2} (X_n - \mu_k)^T \Theta_k (X_n - \mu_k) + \frac{1}{2} \log \det \Theta_k - \frac{DR}{2} \log 2\pi, \quad (5)$$

where μ_k refers to the empirical mean of the k th cluster, and $\det \Theta_k$ denotes the determinant of the matrix Θ_k .

2.3. Cluster assignments and parameter updates

Algorithm 1 Overall steps for STICC

initialize cluster assignments \mathbf{P} and cluster parameters Θ
while not stationarity
 E-step: cluster assignments $\rightarrow \mathbf{P}$
 M-step: parameter updates $\rightarrow \Theta$
endwhile
 return \mathbf{P}, Θ

To solve the problem defined in Equation (3) that involves both combinatorial and continuous optimization, we adopt an approach similar to the expectation-maximization (EM) method that alternates between two steps: cluster assignments and cluster parameters updates. As shown in Algorithm 1, the overall procedure of the STICC algorithm is as follows:

1. **Initialization:** initialize cluster assignments \mathbf{P} and cluster parameters Θ . The former could be randomly initialized, or initialized using other clustering methods, if available, such as K -Means or Gaussian mixture model.
2. **E-step: compute** the set of cluster assignments \mathbf{P} for the geographic objects.
3. **M-step: update cluster** parameters Θ .
4. If the algorithm **converges**, then stop; otherwise, repeat steps (2) and (3).

Further details of the E-step and M-step are described in the following sections.

2.3.1. E-step: cluster assignment with spatial consistency

In this step, we aim to assign each geographic subregion, X_1, \dots, X_N , to one of the K clusters given the fixed cluster parameters $\Theta_i, i = 1, 2, \dots, K$ (i.e. inverse covariances). The number of clusters K should be assigned manually. In practice, K can be inferred by using methods for unsupervised clustering such as silhouette score (Ogbuabor and Ugwoke 2018), elbow method (Syakur *et al.* 2018), information criterion approach (Kodinariya and Makwana 2013), and information-theoretic approach (Sugar and James 2003). Particularly, the optimization problem for assignment of subregions to clusters is derived from the Equation (3) and given by

$$\min_{\mathbf{P}} \sum_{k=1}^K \sum_{X_n \in P_k} \left(\underbrace{-\mathcal{L}(\Theta_k; X_n)}_{\text{log likelihood}} + \underbrace{\beta \mathbb{1}\{X_n^{(1)} \notin P_k\}}_{\text{spatial consistency}} \right). \quad (6)$$

As described in Section 2.2.3, this formulation aims to maximize both the log likelihood and the spatial consistency.

A naive approach to solve the above combinatorial optimization problem is by enumerating all possible assignments of the subregions to the clusters, which, however, quickly becomes infeasible as there are K^N possible combinations, leading to




Figure 4. Problem (6) can be converted to a minimum cost path finding task from subregion 1 to N , where the node cost is the negative log likelihood of that point being assigned to a given cluster, and the edge cost is determined by the function whether the n th geographic object and its nearest neighbor belong to the same cluster; if so, then it is 0, otherwise, a penalty β is added.

exponential increase of running time in the number of geographic objects. To handle this issue, we leverage a dynamic programming approach to solve the problem with a running time $O(KN)$, as illustrated in the Figure 4. Such a problem can be cast into a minimum cost path finding task (Viterbi 1967) from subregion X_1 to X_N . As shown in Figure 4, the node cost refers to the negative log likelihood $-\mathcal{L}(\Theta_k; X_n)$ of a specific cluster that the geographic object is assigned to. We define the edge cost from the node $-\mathcal{L}(\Theta_k; X_n)$ to the destination node $-\mathcal{L}(\Theta_{k'}; X_{n+1})$ to be $h(X_n, X_n^{(1)}, P_k)$ that is decided based on the cluster assignment of the nearest subregion to X_n , and formally defined as

$$h(X_n, X_n^{(1)}, P_k) = \begin{cases} 0, & \text{if } X_n, X_n^{(1)} \in P_k, \\ \beta, & \text{otherwise.} \end{cases} \quad (7)$$

For example, in terms of the subregion X_1 , if its nearest subregion $X_1^{(1)}$ belongs to the same cluster, then there is no additional cost except for the negative log likelihood; otherwise, the penalty β should be added. The pseudocode for cluster assignment is illustrated in Algorithm 2.

Algorithm 2 E-step: Cluster Assignment

Input: $\{X_n\}_{n=1}^N$: a set of N subregions;
 $\{X_n^{(1)}\}_{n=1}^N$: a set of N subregions that indicates the nearest neighbor of X_n ;
 $\mathcal{L}(\Theta_k; X_n)$: negative log likelihood of the subregion X_n that belongs to cluster k , for $k = 1, 2, \dots, K$ and $n = 1, 2, \dots, N$;

$\beta > 0$: penalty coefficient.

initialize FutureCostVals := list of N lists of K zeros.

for $n = N-1, \dots, 1$ **do**

m := index of $X_n^{(1)}$

FutureCost := FutureCostVals[m , :]

LogLikelihoodVals := $(\mathcal{L}(\Theta_1; X_n^{(1)}), \mathcal{L}(\Theta_2; X_n^{(1)}), \dots, \mathcal{L}(\Theta_K; X_n^{(1)}))$

for $k = 1, \dots, K$ **do**

TotalVals = FutureCost + LogLikelihoodVals + β

TotalVals[k] = TotalVals[k] - β

FutureCostVals[n , k] = min(TotalVals)

end for

end for

Path := list of N zeros.

for $n = 1, \dots, N$ **do**

m := index of $X_n^{(1)}$

FutureCost := FutureCostVals[m , :]

LogLikelihoodVals := $(\mathcal{L}(\Theta_1; X_n^{(1)}), \mathcal{L}(\Theta_2; X_n^{(1)}), \dots, \mathcal{L}(\Theta_K; X_n^{(1)}))$

TotalVals = FutureCost + LogLikelihoodVals + β

TotalVals[Path[n]] = TotalVals[Path[n]] - β

end for

return Path

As suggested by the objective function (6), setting β to 0 reduces to independently assigning each subregion X_1, \dots, X_N to a certain cluster based on its log likelihood, which, as explained in Section 1, is not desirable as it is not much different from attribute-based clustering that does not adequately incorporate the geographic information. With the increase of β , nearby subregions are encouraged to be allocated to the same cluster to keep the spatial contiguity. For the extreme case where $\beta \rightarrow \infty$, all subregions will be grouped into the same cluster as the penalty term dominates entirely.

2.3.2. M-step: cluster parameter updates with *Toeplitz graphical lasso*

The other step is the parameter updates given the fixed cluster assignments \mathbf{P} . It aims to update parameters Θ_k of cluster $k = 1, \dots, K$. Here, we assume that the multivariate Gaussian distribution of X_n has zero-mean, and rewrite the log likelihood term in Equation (5) as

$$\sum_{X_n \in P_k} \mathcal{L}(\Theta_k; X_n) = -|P_k|(\log \det \Theta_k + \text{tr}(S_k \Theta_k)) + C, \quad (8)$$

where $|P_k|$ indicates the number of geographic objects assigned to cluster k , S_k denotes the empirical covariance of these objects, tr refers to tracing over the diagonal elements of the matrix, and C is a constant.

To update the parameter Θ_k , the optimization subproblem of the M-step can be written as

$$\min_{\Theta_k \in \mathcal{I}} -\log \det \Theta_k + \text{tr}(S_k \Theta_k) + \frac{1}{|P_k|} \|\lambda \odot \Theta_k\|_{1, \text{off}}. \quad (9)$$

The above problem, termed as the *Toeplitz graphical lasso* (Hallac *et al.* 2017), is convex, as is the case in the typical setting of graphical lasso (Friedman *et al.* 2008).

A straightforward approach to solving this optimization problem is to use the standard coordinate descent method developed by Friedman *et al.* (2008), which, however, may not scale up to large inverse covariances, especially since our EM-like procedure requires solving this problem up to tens or hundreds of times before converging to a stationary solution. Other improved alternatives include the second-order approach with quadratic approximations (Hsieh *et al.* 2014), or the alternating direction method of multipliers (ADMM) (Boyd *et al.* 2010, Scheinberg *et al.* 2010). In this work, we adopt the ADMM method, a distributed algorithm that solves an optimization problem by decomposing it into smaller subproblems, each of which are easier to solve. Such a strategy is also used in the TICC algorithm (Hallac *et al.* 2017) and more technical details can be found there.

3. Experiment

We have implemented the proposed STICC algorithm in Python. The PySal library is used for constructing the spatial relationships, performing the k-nearest neighbor spatial matrix, and constructing the Delaunay triangulation (Rey and Anselin 2010). Following the reproducibility and replicability guidelines (Wilson *et al.* 2020), the code for this project is open-sourced and available on a public repository on GitHub: <https://github.com/GeoDS/STICC>. In the following sections, the developed STICC algorithm is applied in two case studies, a synthetic experiment and a real-world scenario to demonstrate how this method can be used in spatial clustering to provide meaningful insights for repeated geographic pattern discovery.

3.1. Experiment set-up

In this section, we discuss the common settings of the two experiments, including baseline methods and evaluation metrics. Both experiments are performed on a cloud server with Ubuntu 16.04 system and Python 3.8 version.

3.1.1. STICC and baseline methods

There are two groups of clustering methods used in our experiments. The first group includes a series of the proposed STICC clustering algorithms with different parameter settings on R and β , so that we can explore the performance of the algorithm under different conditions. We start with a fixed β and an increasing R value from $R=1$ to $R=4$. When $R=1$, no nearby geographic objects are used for the construction of sub-regions. Then, by picking up the best R , we test different $\beta = 0, 1, 3, 5$, respectively. It should be noted that when $\beta=0$, there is no spatial consistency strategy performed.

The second group contains the following six commonly used unsupervised clustering algorithms to serve as baseline approaches.

1. *K*-Means: The *K*-Means clustering algorithm with Euclidean distance in a multivariate space (including only aspatial attributes in our experiments) (Anderberg 1973).
2. Spatial *K*-Means: It is an adaption of the *K*-Means that treats the coordinates of each object as two attributes together with other aspatial attributes.
3. GMM: The Gaussian mixture model aims at representing the probability distribution of each object for clustering (Rasmussen 1999).
4. CURE: The clustering using representatives model that extends *K*-Means clustering approach and is more robust to outliers, shape and size variances (Guha *et al.* 1998).
5. DBSCAN: The density-based spatial clustering of applications with noise (Ester *et al.* 1996).
6. Spatially Constrained Multivariate Clustering: The spatially constrained multivariate clustering algorithm that uses minimum spanning tree and a strategy termed as SKATER that can identify natural clusters and evidence accumulation to evaluate the probability of cluster membership (Assunção *et al.* 2006).

The first four methods are attribute-based clustering algorithms; DBSCAN algorithm is a density-based clustering approach; and the spatially constrained multivariate clustering approach is a regionalization-based approach. Therefore, the characteristics of the three-category spatial clustering approaches can be depicted comprehensively and compared with our proposed STICC method. All algorithms except GMM and spatially constrained multivariate clustering are performed using a python package *pyclustering* (Novikov 2019); the GMM is performed using the package *sklearn*; the spatially constrained multivariate clustering is executed in ArcGIS Pro 2.8. We use default settings of each algorithm as they mainly require the number of clusters *K* only. A grid-search strategy is performed for DBSCAN to infer optimal results, since this algorithm requires two input parameters, namely the search *radius* and the minimum number of data points *minPts* in each cluster.

3.1.2 Evaluation metrics

To evaluate the performances of all these above-mentioned clustering methods, we use the following three metrics from different aspects: adjusted rand index (ARI), macro-F1, and join count ratio. Generally, clustering is treated as an unsupervised learning process. However, since we also know the ‘ground-truth’ results of the clustering in our case studies, it can be seen as a supervised multi-class classification problem as well. The ARI and macro-F1 are good indicators in measuring the accuracy of clustering results. While the join count ratio can be used to measure the spatial dependence. Given that the STICC method and several baseline methods may acquire the appropriate number of clusters *K*, we assign the ‘true’ number of clusters of each dataset as *K*.

The ARI measures the similarity between two clustering results by counting pairs of samples that are assigned with the same or different clusters (Steinley 2004). A value between 0 and 1 is produced for each round of comparison between clustering results. The higher the value, the more similar the two clustering results.




Figure 5. The generated synthetic dataset with ten regions (R_1, R_2, \dots, R_{10}) that belong to seven clusters.

The macro-F1 score has been widely used in the field of machine learning for measuring algorithm performance (Fujino *et al.* 2008). The precision and recall of each cluster are computed at first to infer the F1 scores. The average of all F1 scores for all clusters are calculated as the macro-F1 score for making comparisons between the STICC method with other baseline models. Such an index achieves the balance between precision and recall. The macro-F1 score ranges from 0 to 1 as well. A value near 1 indicates a better clustering result. To obtain the macro-F1 score, given that we do not know the one-to-one cluster label matching between the ‘ground truth’ and the output clustering results, all potential permutations are enumerated and the macro-F1 score of each item is computed. The permutation with the highest macro-F1 score is used as the final clustering labels.

Researchers have made extensive use of the former two indices for measuring the performance of clustering approaches. Recall that an ideal spatial clustering result may achieve a balance between attribute similarity and spatial contiguity. Hence, a quantitative examination of the degree of spatial contiguity that is preserved in clustering results is necessary. In other words, how many nearby geographic objects tend to belong to the same cluster. Inspired by the join count statistics (Dacey 1965, Cliff and Ord 1973), we measure the proportion of connections where neighboring geographic objects belong to the same cluster to all connections among adjacent geographic objects, which is termed as textitjoin count ratio. For a pair of two neighboring geographic objects, their resulting cluster labels might be the same or different. It should be noted that, for polygons, neighboring geographic objects refer to adjacent polygons or lattices with shared boundaries; while for points, Delaunay triangulation or other types of connectivity among points should be constructed first, and neighboring points are considered as adjacent geographic objects. We adopt the following equations to calculate the join count ratio:

$$J = J_{\text{same}} + J_{\text{diff}} \quad (10)$$

$$J_{\text{ratio}} = \frac{J_{\text{same}}}{J} \quad (11)$$

where J_{same} is the value of join counts that neighboring objects belong to the same cluster, J_{diff} indicates the value of join counts that neighboring objects belong to different clusters.

Table 1. Settings of the generated synthetic dataset with five attributes.

Cluster	Attribute A		Attribute B		Attribute C		Attribute D		Attribute E	
	μ	θ	μ	θ	μ	θ	μ	θ	μ	θ
1	4	1	1	3	80	20	1000	350	999	3
2	5	1	7	3	30	20	900	350	992	3
3	6	1	2	3	20	20	600	350	1005	3
4	1	1	3	3	100	20	700	350	1003	3
5	3	1	6	3	60	20	800	350	999	3
6	7	1	4	3	70	20	400	350	998	3
7	2	1	5	3	40	20	500	350	1008	3

J_{ratio} ranges from 0 to 1, and the higher the proportion of neighboring geographic objects belonging to the same category, the higher spatial dependence of output results.

Such a measure may help us examine the spatial contiguity within each subregion given the spatial dependency of attributes. We acknowledge that there are multiple other measures used for examining spatial contiguity of geographic patterns, such as Moran's I and Geary's C. However, they are inappropriate for nominal data (e.g. clustering labels) in this work.

3.2. Synthetic experiment

3.2.1. Dataset Preparation

We first generate a synthetic point dataset with multiple attributes to test the performance of our proposed STICC method. Inspired by previous works (Estivill-Castro and Lee 2002, Nosovskiy *et al.* 2008, Liu *et al.* 2012, Mai *et al.* 2018), the spatial distribution of points are generated as shown in the Figure 5 with the following characteristics. According to the figure, regions R_1 and R_4 , regions R_2 and R_{10} , and regions R_3 and R_9 , belong to the same clusters, respectively, but they may locate in different positions and are not spatially adjacent to each other (e.g. R_3 and R_9 , R_2 and R_{10}). These regions have diverse densities of points. For example, regions 4 and 8 have uniform point densities, while point densities fluctuate in regions 5 and 6; densities of points in adjacent regions may be similar (e.g. regions 7 and 9) or different (such as regions 7 and 8). Also, the minimum distances between two nearby regions can be different. For instance, region 5 is closely connected (i.e. 'touch') with region 6 while is not directly adjacent to region 2.

Though there are ten regions of points in total, they only belong to seven predefined clusters according to their multiple attribute similarity. For each point, we generate five attributes as its multi-dimensional features. The attribute values of each point in a specific cluster are randomly assigned in a certain range, and such a range follows normal distributions with distinct settings of mean μ and standard error θ . After randomly generating these parameters for multiple times, similar conclusions can be obtained from our further experiments. Hence, we only show one group of parameters as an example in Table 1 for demonstrating the experiment settings. Table 1 provides the means μ and standard errors θ of the normal distribution for the attribute value range of objects in each cluster. Attribute values may overlap in different clusters due to the settings of θ . Given that different attributes of one geographic object may have diverse value ranges in the real world (e.g. the average temperature and the

Table 2. Results of different clustering approaches on synthetic dataset (number of clusters, $K=7$).

	Cluster		Adjusted rand index	Macro-F1	Join count ratio
	R	β			
STICC	1	3	0.894	0.954	0.851
	2	3	0.931	0.973	0.878
	3	3	0.960	0.984	0.901
	4	3	0.574	0.550	0.822
	3	0	0.947	0.977	0.888
	3	1	0.952	0.981	0.896
	3	5	0.818	0.771	0.886
Baseline clustering methods	<i>K</i> -Means		0.799	0.735	0.544
	CURE		0.0006	0.053	–
	Spatial <i>K</i> -Means		0.830	0.744	0.881
	GMM		0.703	0.850	0.685
	DBSCAN (<i>radius</i> = 1250, <i>minPts</i> = 25)		0.327	–	0.933
	Spatially constrained multivariate clustering		0.629	0.546	0.936

population of a place are unrelated), we assign disparate settings of normal distributions to create these attributes. As suggested in Table 1, the attribute A ranges from 1 to 7 while the attribute D varies from 400 to 1000.

3.2.2. Clustering results

The comparison results of different clustering methods are shown in Table 2. Several selected clustering results are displayed in Figure 6. To guarantee that model performances are compared under same conditions, we generate identical synthetic data for each of the methods.

3.2.2.1. Accuracy. In accordance with the Table 2, our proposed STICC method significantly outperforms all baseline clustering approaches in identifying such dispersed distributed similar subregions regarding both ARIs and macro-F1 scores. The STICC method with different parameters achieves very high accuracy in most cases. The ARIs of all experiments but one are higher than 0.81, and the macro-F1 scores in five out of seven experiments are over 0.95 with STICC. The model with parameters $R=3$ and $\beta=3$ performs the best. Its macro-F1 score reaches 0.984. Figure 6(A) shows the results of STICC ($R=3, \beta=3$). As shown in the figure, all clusters are correctly identified and the majority of points are classified to the correct group. Repeated clusters (e.g. R_3 and R_9 , R_2 and R_{10}), even if located in different positions, are detected accurately. Only a small proportion of points are labeled incorrectly (as indicated by the macro-F1 score, more than 98% points are classified correctly).

As for the baseline approaches, when setting the number of clusters $K=7$, spatial *K*-Means has the best ARI of 0.83 (0.13 less than the best STICC), and GMM has the highest macro-F1 of 0.850 (0.134 less than the best STICC). Figure 6(B–D) depict the results of those baseline clustering approaches. Results of *K*-Means show that most clusters are identified (Figure 6(B)). However, regions 7 and 8 are grouped into the same cluster incorrectly; regions 2 and 10, which belong to the same group, are misclassified. According to Figure 6(C), coordinates of points are treated as two attributes when executing the spatial *K*-Means clustering, nearby points are encouraged to be

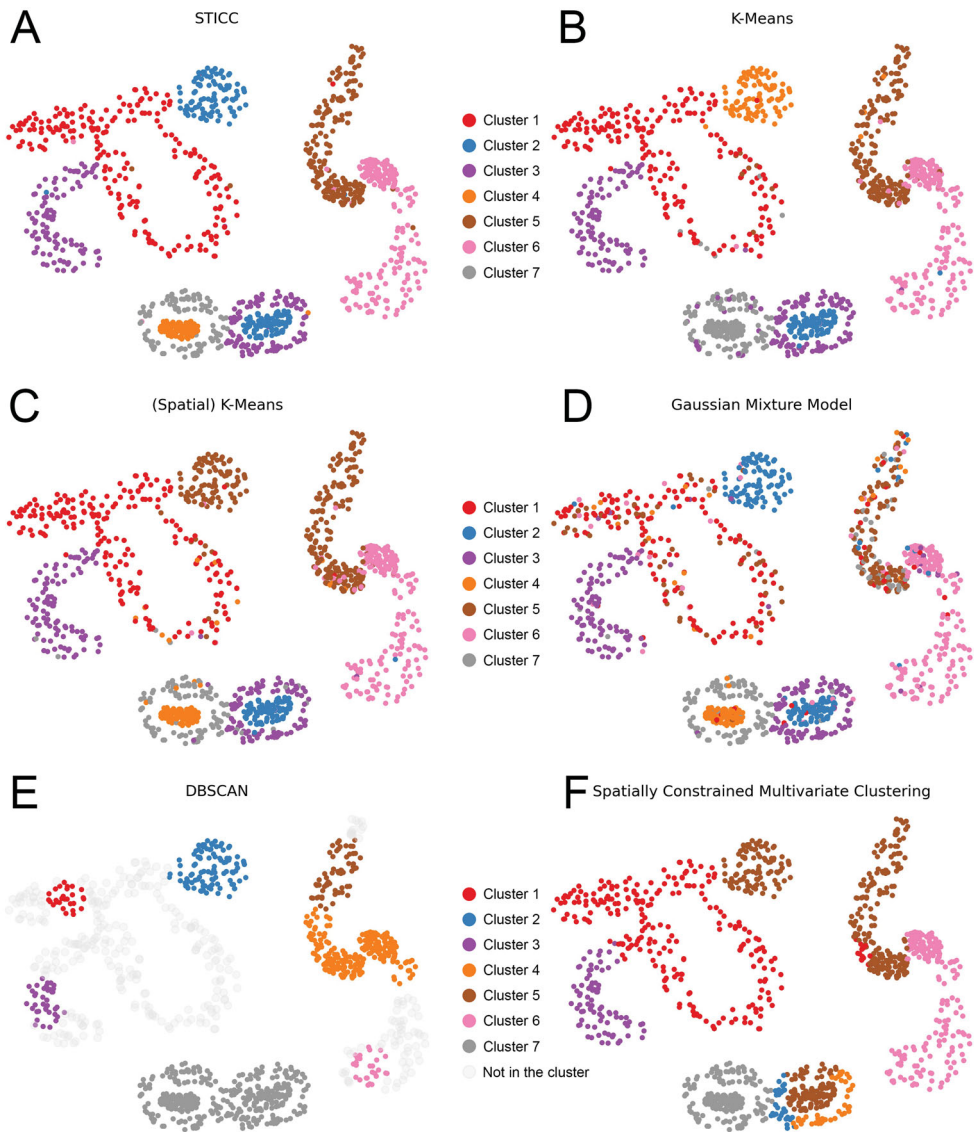


Figure 6. Results of the synthetic dataset (seven clusters). (A) the proposed STICC method (with $r=3$ and $\beta=3$); (B) K-Means; (C) spatial K-Means; (D) Gaussian mixture model; (E) DBSCAN (with $radius = 1250$ and $minPts = 25$); (F) spatially constrained multivariate clustering.

grouped into the same cluster. In the synthetic dataset, such a method performs worse than K-Means as regions 2 and 5 are misclassified into the same group. GMM, as shown in Figure 6(D), generally detects all major clusters by attributes correctly. However, in comparison with the STICC method, the clustered points are in a more disordered way without good spatial contiguity. In terms of density-based clustering algorithm (DBSCAN), according to Figure 6(E), it can only discover densely-located point clusters based on locations but not attributes. Therefore, many points are not assigned into any clusters but as noise points, resulting in a relatively low ARI and a failure in getting a macro-F1 score as more than 7 clusters are obtained. For the

spatially constrained multivariate clustering using ArcGIS Pro, such a method can detect certain clusters (Figure 6(F)). When performing the clustering algorithms, the Delaunay triangulations are constructed among points. Only topologically nearby points are assigned into the same group and no repeated patterns are discovered.

3.2.2.1. Spatial contiguity. According to Table 2, the spatial contiguity (measured by the join count ratio) is well maintained by our proposed STICC method. When $R = 2, \beta = 3$, and $R = 3, \beta = 0, 1, 3, 5$, the join count ratio of all these five experiments are higher than that of the K-Means and GMM approaches. It should be noted that a higher join count ratio is not necessarily to a higher ARI and macro-F1, as they measure different aspects of the model performance. For example, if we only compare the K-Means algorithm and the spatially constrained multivariate clustering approach, the latter one has a higher join count ratio value but with lower ARI and macro-F1 score. It makes sense as illustrated in Figure 6(B,C). K-Means can identify correct clusters to a certain degree, though there are some noise points that are misclassified; while spatial K-Means incorrectly arrange the R_2 and R_5 together, as coordinates of points in these two regions are relatively similar. In addition, both DBSCAN and the spatially constrained multivariate clustering approach have relatively high join count ratio, which indicates that the spatial contiguity is well preserved. However, the ARI and macro-F1 scores, which measure classification accuracy, are relatively low for these two methods. This suggests that the join count ratio may only serve as a complement evaluation metric of ARI and marco-F1 scores from different perspectives.

3.2.2.2. Parameters of STICC. We also try multiple combinations of R, β values to see how these parameters affect clustering results based on our synthetic dataset. With the increase of R and fixing the β , in which more nearby points are used for constructing the subregions, the accuracy of clustering approaches improves at first, and then drops when $R = 4$. When we fix the R nearest neighbors and change β , we can observe similar patterns (Table 2). The three indices increase at first and then drop when $\beta = 5$. A larger β forces more nearby objects to be assigned into the same group.

In summary, the STICC approaches perform better than the other three groups of methods significantly in discovering repeated patterns of points with spatial contiguity maintained, and achieve the balance between spatial and aspatial attributes of geographic objects with appropriate parameters.

3.3. Check-in point classification

We then evaluate the STICC method as well as other baseline approaches on another real-world application: social media check-in point classification. The goal of this task is to extract semantic information such as place types from geographic coordinates and time stamps only. It is very useful for multiple applications such as trajectory privacy protection (Rao *et al.* 2020), mobility pattern discovery (Soliman *et al.* 2017), and community detection (Zhao *et al.* 2016).

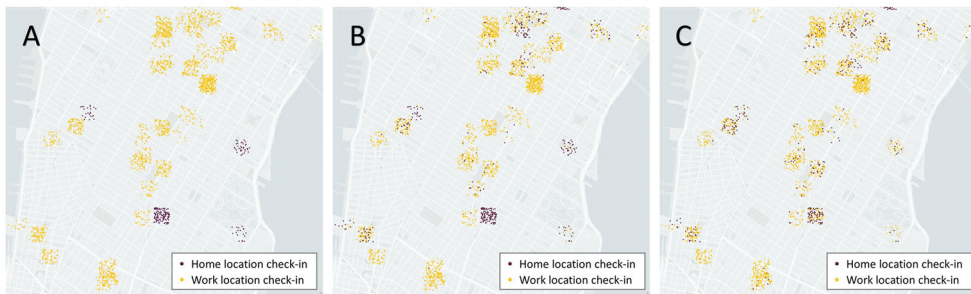


Figure 7. Results for home/work classification. (a) ground truth; (b) the proposed STICC method; (c) K-Means.

3.3.1. Dataset preparation

Social media check-in datasets have been widely used for understanding and analyzing users' spatio-temporal activity patterns in location-based services (LBS). Users visit a place (usually referred to as a point of interest, POI), check in at that location, and post geotagged content. Our testing data are from Yang *et al.* (2014), which contain a Foursquare check-in dataset in New York City processed based on global social media check-in records. In this dataset, 118,316 social media check-in records from April 2012 to September 2013 were collected from Foursquare, containing 3628 unique POIs and 500 users. For each check-in record, the time stamp (including time of day and day of week), GPS coordinates of the POI, and its semantics represented by the POI categories are attached. We use the POI category provided in the original dataset as the ground truth in this application. This task aims at identifying POI categories based on temporal information of check-in records. Given that some POIs have very limited visits, only those with at least 10 check-in points are kept. We then extract days of the week $w \in (1, 7)$, and the hour $h \in (0, 23)$ information from time stamps as two attributes of each point. Given that check-in points attached to one POI have the same GPS coordinates, we add a random noise (following a uniform distribution) of every point so that each of them is shifted to locate at a different position while still inside the 100-m radius of the center POI. By doing so, the task becomes more challenging for clustering algorithms to correctly differentiate check-in points in mixed groups of sampling points. After that, we test the STICC method as well as other baseline approaches to classify the category of check-in points on two tasks, home/work, and home/work/gym identification. Again, the process is still an unsupervised learning problem without labels and the identification of place types requires temporal information (e.g. daytime vs. night time) after clustering. Existing studies have shown that the day of the week and the hour play important roles in identifying various place types (Wu *et al.* 2014, Rao *et al.* 2020). With only two attributes taken into account, we explore to what degree correct types of places can be inferred. Only attribute-based clustering methods are used for comparisons with STICC. Density-based clustering methods can detect densely located point clusters. However, it does not identify the semantics of each cluster. Regionalization-based clustering methods may only partition the study area into two spatially adjacent groups. Hence, the latter two methods are not appropriate for this application.

Table 3. Results of different approaches for home and work classification (number of clusters, $K = 2$).

	Cluster		Adjusted rand index	Macro-F1	Join count ratio
	R	β			
STICC	1	3	0.390	0.806	0.829
	2	3	0.355	0.792	0.804
	3	3	0.433	0.823	0.834
	4	3	0.495	0.844	0.860
	4	0	0.445	0.823	0.822
	4	1	0.464	0.834	0.841
	4	5	0.514	0.850	0.871
	Traditional clustering	K-Means		0.085	0.321
CURE		0.015	0.578	0.700	
Spatial-Kmeans		0.080	0.384	0.492	
GMM		0.023	0.587	0.690	

3.3.2. Clustering accuracy for home/work identification

Home and work location detection is one of the fundamental tasks for LBS in practice. Here, we use different clustering algorithms and perform a binary ‘classification’ ($K = 2$) to infer users’ home and work locations from their check-in data. The distributions of users’ check-in points are plotted in Figure 7(A) with yellow dots indicate home locations and purple points refer to work locations. The clustering results of home/work identification of several clustering methods are demonstrated in Table 3. It can be inferred that the proposed STICC algorithm performs better than other baseline approaches using the three quality measurements.

3.3.2.1. Accuracy. According to Table 3, the ARI values of the STICC with different parameter settings are all over 0.35, and the macro-F1 scores are all higher than 0.79. When $R = 4, \beta = 5$, the STICC has the highest ARI, macro-F1 score, and join count ratio. In comparison, the K-Means is the best model with the highest ARI of 0.085, and GMM has the highest macro-F1 score of 0.587, both are far below the results from the STICC algorithm. Three example maps are illustrated in Figure 7 to show the spatial distributions of check-in ground truth, STICC, and K-Means clustering results, respectively. According to these maps, our developed STICC algorithm identifies home/work locations of most check-in points more correctly than the K-means, although there exist some mixed home/work clusters in the northern part on the map.

3.3.2.2. Spatial contiguity. The join count ratio also illustrates that the STICC method maintains the spatial contiguity relatively well. According to Table 3, join count ratios are all greater than 0.80. Also, as shown in Figure 7(B), nearby check-in points are mostly grouped into the same cluster, which also suggests that spatial contiguity is well-preserved. In comparison, attribute-based algorithms cannot keep the spatial contiguity well with relatively low join count ratio values (lower than 0.70).

3.3.2.3. Parameters of STICC. Since the model performs best when $R = 4$, multiple β values are input into the STICC method to examine the resulting clustering maps for the POI classification task. As shown in Table 3, as the number of objects in each sub-region increases, the accuracy and join count ratio of the STICC method also show

Table 4. Results of different approaches for home/work/gym classification (number of clusters, $K = 3$).

	Cluster		Adjusted rand index	Macro-F1	Join count ratio
	R	β			
STICC	1	3	0.289	0.476	0.700
	2	3	0.204	0.482	0.647
	3	3	0.269	0.500	0.672
	4	3	0.298	0.508	0.700
	4	0	0.251	0.495	0.641
	4	1	0.273	0.502	0.669
	4	5	0.335	0.510	0.712
Traditional clustering	K-Means		0.041	0.352	0.675
	CURE		0.077	0.294	0.597
	Spatial-Kmeans		0.080	0.397	0.670
	GMM		0.065	0.416	0.603

increasing trends. Also, if β increases, i.e. more nearby geographic objects are pushed to be in the same group, all three quality measurements increase as well.

3.3.3. Clustering Accuracy for home/work/gym identification

We then conduct a follow-up experiment to perform multi-class ‘classification’ ($K = 3$) for grouping and identifying home, work, and gym locations. Since only the temporal patterns (hour and days of the week information) of check-in points are considered, it is an obstacle for most clustering algorithms even just extending the experiment from binary classification to three-class classification. Table 4 shows the three quality measurements of such a task. As mentioned above, all potential cluster label matching permutations are enumerated and the best results are reported here. Though the performances of all methods drop substantially, our proposed STICC algorithm continues yielding best results.

Generally speaking, similar conclusions can be obtained in comparison with the home/work classification task. As for ARI, the proposed STICC yields an ARI that is at least 0.12 higher than any other methods. For macro-F1 scores, the STICC achieves the highest when $R = 4, \beta = 5$ which is higher than GMM (0.416) that performs the best among all other clustering approaches. The resulting join count ratios of most STICC algorithms are higher than 0.64 with the highest being 0.712, while K-Means has the highest join count ratio with 0.675 among traditional clustering methods. It is worth noting that when $R = 4, \beta = 0$, the join count ratio is comparatively low due to the lack of a spatial consistency strategy that forces nearby points to be in the same group.

Overall, both home/work and home/work/gym identification practices illustrate that our proposed STICC algorithm shows promising results in POI category identification problems inferred solely from hourly and days of the week. Our results also suggest that it is necessary to consider nearby geographic observations in improving the model performance.

4. Discussions

4.1. Influences of parameters

The developed STICC algorithm requires four input hyperparameters K , R , β , and λ , which are manually determined. To get optimal results, fine tuning is required when




Figure 8. Changes of metrics for the synthetic dataset with different R .

executing the algorithm. In this section, we briefly discuss the influence of these different parameters on the algorithm performance and hope to provide basic guidance in selecting hyperparameters.

K refers to the number of clusters into which the data points expect to be grouped. If there are some labeled ground-truth data, such data can be used for helping determine the K value, as illustrated in the experiments shown in Section 3. However, clustering algorithms are usually employed in solving unsupervised classification problems where no ground-truth data are provided. Researchers may either need to manually select the appropriate K relying on their experiences or based on the characteristics of datasets, or need to use silhouette score (Ogbuabor and Ugwoke 2018), elbow method (Syakur *et al.* 2018), information criterion (Kodinariya and Makwana 2013), or information-theoretic (Sugar and James 2003) approaches to help determine the appropriate K .

R specifies the radius size of subregions, that is the number of geographic objects in each constructed subregion in our method. It also serves as the key parameter in integrating spatial context into the clustering algorithm. The larger the R , the more neighbors are taken into account in clustering; and if $R = 1$, only the center point itself is used. The hypothesis for constructing subregions is that spatial dependencies among geographic objects within subregions are the same for a specific cluster and might be different across clusters. Hence, choosing an appropriate R value depends on the nature of the phenomenon being studied and the dataset. For example, $R = 5$ or 9 might be proper parameters for raster data (e.g. remote sensing images) considering the adjacency characteristics of image pixels (such as rook, bishop, and queen cases). In general, a large R might be suitable when objects are highly relevant to their neighbors. With a large R , more geographic objects are taken into consideration when the STICC algorithm is executed, which also increases the spatial variances. On the




Figure 9. Results of the synthetic dataset using STICC with different β . (A) $\beta = 0$, (B) $\beta = 3$, (C) $\beta = 18$.

contrary, a small R may be used for datasets with high spatial heterogeneity, in which case fewer geographic objects should be taken into account. Here, we take the synthetic dataset as an example to show how results change with the increasing R . We perform our STICC algorithm with different R ranging from 1 to 10. As illustrated in Figure 8, both ARI and macro-F1 scores increase and reach the peak at $R = 2$ or $R = 3$, which means that taking several nearby geographic objects into account helps illustrate spatial dependencies within subregions; the two metrics then drop when R is greater than 4. Thus, taking more nearby geographic objects into account may provide more complex spatial context information that could ‘confuse’ the model. In terms of the join count ratio, it is relatively stable when R is less than 8, indicates that spatial contiguity of the observations is well preserved, and decreases when R is greater than 9. Note that this pattern is only observed based on our synthetic dataset.

The β controls the penalties of the spatial consistency strategy. The larger the β , the more nearby points are grouped into the same cluster; while if $\beta = 0$, no such strategy is carried out. An optimal β might be inferred using parameter tuning. Such a parameter impacts on the spatial contiguity maintained in the results. Using the abovementioned synthetic dataset (in Figure 5), we provide an example here to show how clustering results change with different β . As illustrated in Figure 9, three sub-figures illustrate clustering results with $\beta = 0, 3, 18$, respectively. When β is 0, no spatial consistency strategy is performed, and the points are relatively messy as points with different classes mixed up. While with a relatively large $\beta = 18$, nearby geographic objects are grouped into the same cluster, such as regions 7 and 9 (in purple) at the bottom, though they should belong to different clusters based on their attributes.

Recall that the hyperparameter λ controls the level of sparsity on the inverse covariance matrices and can be used to prevent overfitting issues. Although the inverse covariance matrix consists of a $DR \times DR$ matrix in its general form, we restrict all the entries to a single value to save the hyperparameter search cost. In practice it could be picked via cross-validation.

4.2. Clustering result interpretation

In this section, we discuss an approach for interpreting clustering characteristics. Since the structure of each cluster outputted from STICC is represented as a multilayer MRF

Table 5. Betweenness centrality of attributes in different clusters of the synthetic dataset.

	Attribute A	Attribute B	Attribute C	Attribute D	Attribute E
Cluster 1	0.00	0.00	0.50	0.83	0.00
Cluster 2	0.00	0.00	0.50	0.83	0.25
Cluster 3	0.83	0.00	0.08	0.83	0.58
Cluster 4	0.42	0.00	0.25	0.58	0.92
Cluster 5	0.17	1.00	0.00	0.42	0.67
Cluster 6	0.83	0.42	0.17	0.58	0.00
Cluster 7	0.00	0.17	0.92	0.33	0.83

The higher the betweenness centrality, the more 'important' the variable is in determining the cluster.

network, network analysis approaches can be used for evaluating the properties of each cluster. Betweenness centrality that assesses the number of shortest paths passing through the node has been frequently used as a measure of centrality in networks (Crucitti *et al.* 2006). It emphasizes the importance of each node. Taking the synthetic dataset as an example, the betweenness centrality score of each node in the output MRF network is computed and shown in Table 5. According to the table, the betweenness centrality of each attribute in each of the seven clusters is different. It can be inferred that attributes C and D play important roles in cluster 1, while attributes A, D and E are relatively important for cluster 3. Though for the synthetic dataset, no physical meanings are attached to each variable, using network analysis to examine MRF networks provides possible ways in characterizing properties of each cluster. In addition, not only the betweenness centrality might be used for cluster interpretation, but also other network measures, such as degree centrality and closeness centrality, might be utilized for different purposes.

4.3. Implications for GeoAI and spatially explicit model

The proposed STICC serves as an example of spatially explicit artificial intelligence techniques for geographic knowledge discovery (Goodchild 2001, Janowicz *et al.* 2020). It shows how spatial thinking can be integrated into artificial intelligence, especially for the development of spatially explicit machine learning algorithms. We made efforts to conceptualize 'space' from the following two aspects. On the one hand, by performing k-nearest neighbor analysis, a subregion is designed for each geographic object to model spatial dependencies among attributes using graph-based MRF. On the other hand, a spatial consistency strategy is utilized by pairing the cluster of each geographic object to its nearest neighbor. Such a strategy, as illustrated by the experiment results, is helpful in maintaining spatial contiguity. If we take steps along these paths, more spatial and aspatial perspectives might be integrated into this algorithm. For instance, in addition to *distant* and *near*, other spatial relationships such as *direction* and *scale* can be modeled into the algorithm (Mai *et al.* 2018, Zhu *et al.* 2019). Spatial matrices such as Gaussian kernel, fixed bandwidth, and distance lags that express different spatial proxy relationships can be embedded into the STICC method as well (Yan *et al.* 2017). In addition, recent studies regarding spatial clustering approaches mainly focus on two directions, either modifying existing algorithms to achieve better performances in classic spatial clustering tasks (Liu *et al.* 2019, Aydin *et al.* 2021), or developing domain-specific algorithms for a specific field such as

cartography (Wolf 2021), human mobility (Liu *et al.* 2021), and geodemography (Grekousis 2021), while limited attention was paid on the proposed RGPD problem. We believe that this work is just a beginning for solving the RGPD problem. There are abundant spatial concepts that can be incorporated into the proposed algorithm so as to promote the state-of-the-art methods in GeoAI.

4.4. Limitations and future work

Although our proposed STICC method has shown promising results and outperformed other clustering methods in both case studies for solving the RGPD problem, we acknowledge the limitations of the proposed method, on which we discuss here. As shown in Figure 6, there are still some outliers in the results as several geographic objects are not labeled correctly. There are two potential reasons: one is that the selected parameters may not be appropriate; the other is that the algorithm may converge early before reaching the global optimum. However, since all clustering approaches have pros and cons, our proposed STICC approach has potentials to solve a series of spatial clustering problems.

In addition, there are two potential directions that are worth exploration in the future work. First, when constructing subregions for each geographic object, we only employ k-nearest neighbors in this research. In fact, more approaches used for defining spatial adjacency relationships might be adapted in designing subregions. For example, a pre-defined distance radius can be used for the definition of subregions, i.e. geographic objects within a specific distance threshold are considered as neighborhoods of the center geographic object for constructing subregions. Under such a circumstance, spatial weights may be further assigned to nearby geographic objects according to inverse distance to construct a weighted neighborhood. Second, given that the repeated geographic patterns are common, we believe that the proposed STICC might be appropriate for a variety of domain applications, including but not limited to climate type classification, semantic trajectory clustering, urban land use detection, remote sensing image segmentation, etc. More experiments using this algorithm and its variations will be conducted in our future work.

5. Conclusion

In this paper, we develop a novel method to discover repeated clusters of multivariate geographic objects considering spatial contiguity. The proposed STICC approach takes nearby geographic objects into account to construct subregions rather than treating each object in isolation. Markov random fields are used for representing dependencies among attributes inside subregions. Then, a spatial consistency strategy is used for forcing nearby geographic objects to be assigned into the same group. An adapted E-M strategy is used for cluster assignment and parameter updates. Two case studies including a synthetic dataset and a set of POI category identification tasks prove that our proposed method has yielded good results in solving the RGPD problem in comparison with other attribute-based, density-based, or regionalization-based clustering methods. It could discover repeated clusters accurately while maintain spatial

contiguity at the same time, which is more in line with the natural distribution pattern of geographical objects. Recall the [Figure 1](#), the STICC method achieves the balance between repeated patterns and spatial contiguity. This new spatial clustering approach can support many other applications in geography, remote sensing, transportation, urban planning, and social sciences.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data and codes availability statement

The data and codes that support the findings of this study are available in figshare.com with the link <https://doi.org/10.6084/m9.figshare.15170898.v1> and on the Github repository <https://github.com/GeoDS/STICC>.

Funding

Yuhao Kang acknowledges the support by the Trewartha Research Award, Department of the Geography, University of Wisconsin-Madison. Song Gao and Jinmeng Rao acknowledge the support by the American Family Insurance Data Science Institute at the University of Wisconsin-Madison and the National Science Foundation funded AI institute [Grant No. 2112606] for Intelligent Cyberinfrastructure with Computational Learning in the Environment (ICICLE). Fan Zhang would like to thank the support by the National Natural Science Foundation of China under Grant 41901321. Ignavier Ng would like to acknowledge the support by the National Institutes of Health (NIH) under Contract R01HL159805. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funders.

Notes on contributors

Yuhao Kang is a Ph.D. student in GIScience at the Department of Geography, University of Wisconsin-Madison. He holds a B.S. degree in Geographic Information Science at Wuhan University. His main research interests include Place-Based GIS, GeoAI, and Social Sensing.

Kunlin Wu is a graduate student in the School of Resource and Environmental Sciences at Wuhan University. His research interests include remote sensing, spatial data mining and auralization of spatial data.

Song Gao is an Assistant Professor in GIScience at the Department of Geography, University of Wisconsin-Madison. He holds a Ph.D. in Geography at the University of California, Santa Barbara. His main research interests include Place-Based GIS, Geospatial Data Science, and GeoAI approaches to Human Mobility and Social Sensing.

Ignavier Ng is a Ph.D. student at Carnegie Mellon University. His research interests include causal discovery and machine learning.

Jinmeng Rao is a Ph.D. student at the Department of Geography, University of Wisconsin-Madison. His research interests include GeoAI, Privacy-Preserving AI, and Location Privacy.

Shan Ye is a Ph.D. candidate at the Department of Geoscience, University of Wisconsin-Madison. His research focuses on quantitative stratigraphy and paleoclimate data science.

Fan Zhang is a Senior Research Associate at Senseable City Lab, Massachusetts Institute of Technology. His research interests include Urban Data Science, Visual Intelligence, GeoAI, and Social Sensing.

Teng Fei is currently an Associate Professor in the School of Resource and Environmental Sciences at Wuhan University. His research focuses on remote sensing, urban data analysis, social sensing and ecological modelling.

ORCID

Yuhao Kang  <http://orcid.org/0000-0003-3810-9450>

Song Gao  <http://orcid.org/0000-0003-4359-6302>

Fan Zhang  <http://orcid.org/0000-0002-3643-018X>

Teng Fei  <http://orcid.org/0000-0002-3415-1654>

References

- Akaike, H., 1973. Block toeplitz matrix inversion. *SIAM Journal on Applied Mathematics*, 24 (2), 234–241.
- Aldstadt, J., 2010. Spatial clustering. In: *Handbook of applied spatial analysis*. Berlin: Springer, 279–300.
- Anderberg, M.R., 1973. *Cluster analysis for applications*. Vol. 19. New York: Academic Press.
- Ankerst, M., et al., 1999. Optics: Ordering points to identify the clustering structure. *ACM SIGMOD Record*, 28 (2), 49–60.
- Aragam, B., and Zhou, Q., 2015. Concave penalized estimation of sparse gaussian Bayesian networks. *Journal of Machine Learning Research*, 16 (1), 2273–2328.
- Assunção, R.M., et al., 2006. Efficient regionalization techniques for socio-economic geographical units using minimum spanning trees. *International Journal of Geographical Information Science*, 20 (7), 797–811.
- Aydin, O., et al., 2021. A quantitative comparison of regionalization methods. *International Journal of Geographical Information Science*, 35 (11), 2287–2315.
- Aydin, O., et al., 2018. Skater-con: unsupervised regionalization via stochastic tree partitioning within a consensus framework using random spanning trees. In: *Proceedings of the 2nd ACM SIGSPATIAL international workshop on AI for geographic knowledge discovery*, 33–42. New York, NY: ACM.
- Baçaõ, F., Lobo, V., and Painho, M., 2005. Self-organizing maps as substitutes for k-means clustering. In: *International conference on computational science*. Berlin, Heidelberg: Springer, 476–483.
- Boyd, S., et al., 2010. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3 (1), 1–122.
- Candes, E.J., and Tao, T., 2005. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51 (12), 4203–4215.
- Chen, Y., et al., 2018. Hispatialcluster: a novel high-performance software tool for clustering massive spatial points. *Transactions in GIS*, 22 (5), 1275–1298.
- Cliff, A.D., and Ord, J.K., 1973. *Spatial autocorrelation*. London: Pion.
- Crucitti, P., Latora, V., and Porta, S., 2006. Centrality measures in spatial networks of urban streets. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics*, 73 (3 Pt 2), 036125.
- Dacey, M.F., 1965. *A review on measures of contiguity for two and k-color maps*. Technical report. Evanston, IL: Northwestern University.

- Deng, M., et al., 2011. An adaptive spatial clustering algorithm based on delaunay triangulation. *Computers, Environment and Urban Systems*, 35 (4), 320–332.
- Deng, M., et al., 2019. A density-based approach for detecting network-constrained clusters in spatial point events. *International Journal of Geographical Information Science*, 33 (3), 466–488.
- Donoho, D.L., 2006. Compressed sensing. *IEEE Transactions on Information Theory*, 52 (4), 1289–1306.
- Duque, J.C., Church, R.L., and Middleton, R.S., 2011. The p-regions problem. *Geographical Analysis*, 43 (1), 104–126.
- Duque, J.C., Ramos, R., and Suriñach, J., 2007. Supervised regionalization methods: A survey. *International Regional Science Review*, 30 (3), 195–220.
- Esri 2021. *How spatially constrained multivariate clustering works*. Available from: <https://pro.arcgis.com/en/pro-app/latest/tool-reference/spatial-statistics/how-spatially-constrained-multivariate-clustering-works.htm> [Accessed 24 June 2021].
- Ester, M., et al., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. *International Conference on Knowledge Discovery and Data Mining*, 96, 226–231.
- Estivill-Castro, V., and Lee, I., 2002. Argument free clustering for large spatial point-data sets via boundary extraction from delaunay diagram. *Computers, Environment and Urban Systems*, 26 (4), 315–334.
- Friedman, J., Hastie, T., and Tibshirani, R., 2008. Sparse inverse covariance estimation with the graphical Lasso. *Biostatistics (Oxford, England)*, 9 (3), 432–441.
- Fujino, A., Isozaki, H., and Suzuki, J., 2008. Multi-label text categorization with model combination based on f1-score maximization. In: *Proceedings of the third international joint conference on natural language processing: volume-II*. Hyderabad, India: IJCNLP.
- Gao, S., Janowicz, K., and Couclelis, H., 2017. Extracting urban functional regions from points of interest and human activities on location-based social networks. *Transactions in GIS*, 21 (3), 446–467.
- Goodchild, M., 2001. Issues in spatially explicit modeling. In: *Agent-based models of land-use and land-cover change*, October 4–7, 13–17. Irvine, CA: NCGIA.
- Goodchild, M.F., 2004. Giscience, geography, form, and process. *Annals of the Association of American Geographers*, 94 (4), 709–714.
- Grekousis, G., 2021. Local fuzzy geographically weighted clustering: a new method for geodemographic segmentation. *International Journal of Geographical Information Science*, 35 (1), 152–174.
- Guha, S., Rastogi, R., and Shim, K., 1998. Cure: an efficient clustering algorithm for large databases. *ACM SIGMOD Record*, 27 (2), 73–84.
- Hallac, D., et al., 2017. Toeplitz inverse covariance-based clustering of multivariate time series data. In: *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 215–223. New York, NY: ACM.
- Hastie, T., Tibshirani, R., and Wainwright, M., 2015. *Statistical learning with sparsity: the Lasso and generalizations*. Boca Raton, FL: Chapman and Hall/CRC.
- Henriques, R., Bação, F., and Lobo, V., 2009. Geosom suite: a tool for spatial clustering. In: *International conference on computational science and its applications*. Springer, 453–466.
- Hinneburg, A., and Keim, D.A., 1998. An efficient approach to clustering in large multimedia databases with noise. *KDD*, 98, 58–65.
- Hsieh, C.-J., et al., 2014. QUIC: quadratic approximation for sparse inverse covariance estimation. *Journal of Machine Learning Research*, 15 (83), 2911–2947.
- Hu, Y., et al., 2015. Extracting and understanding urban areas of interest using geotagged photos. *Computers, Environment and Urban Systems*, 54, 240–254.
- Janowicz, K., et al., 2020. GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, 34 (4), 625–636.
- Kang, Y., et al., 2019. Extracting human emotions at different places based on facial expressions and spatial clustering analysis. *Transactions in GIS*, 23 (3), 450–480.

- Kodinariya, T.M., and Makwana, P.R., 2013. Review on determining number of cluster in k-means clustering. *International Journal*, 1 (6), 90–95.
- Koller, D., and Friedman, N., 2009. *Probabilistic graphical models: principles and techniques*. Cambridge, MA: MIT Press.
- Liqiang, Z., et al., 2013. A spatial cognition-based urban building clustering approach and its applications. *International Journal of Geographical Information Science*, 27 (4), 721–740.
- Liu, K., et al., 2020. Investigating urban metro stations as cognitive places in cities using points of interest. *Cities*, 97, 102561.
- Liu, Q., et al., 2012. A density-based spatial clustering algorithm considering both spatial proximity and attribute similarity. *Computers & Geosciences*, 46, 296–309.
- Liu, Q., et al., 2021. SNN_flow: a shared nearest-neighbor-based clustering method for inhomogeneous origin-destination flows. *International Journal of Geographical Information Science*, 36 (2), 1–27.
- Liu, X., Huang, Q., and Gao, S., 2019. Exploring the uncertainty of activity zone detection using digital footprints with multi-scaled dbSCAN. *International Journal of Geographical Information Science*, 33 (6), 1196–1223.
- MacQueen, J., et al., 1967. Some methods for classification and analysis of multivariate observations. *Berkeley Symposium on Mathematical Statistics and Probability*, 1, 281–297.
- Mai, G., et al., 2018. Adcn: An anisotropic density-based clustering algorithm for discovering spatial point patterns with noise. *Transactions in GIS*, 22 (1), 348–369.
- Miller, H.J., and Han, J., 2009. *Geographic data mining and knowledge discovery*. New York: CRC Press.
- Moayedi, A., Abbaspour, R.A., and Chehreghan, A., 2019. An evaluation of the efficiency of similarity functions in density-based clustering of spatial trajectories. *Annals of GIS*, 25 (4), 313–327.
- Murray, A.T., and Estivill-Castro, V., 1998. Cluster discovery techniques for exploratory spatial data analysis. *International Journal of Geographical Information Science*, 12 (5), 431–443.
- Murray, A.T., and Shyy, T.-K., 2000. Integrating attribute and space characteristics in choropleth display and spatial data mining. *International Journal of Geographical Information Science*, 14 (7), 649–667.
- Ng, I., Ghassami, A., and Zhang, K., 2020. On the role of sparsity and DAG constraints for learning linear DAGs. In: *Advances in Neural Information Processing Systems*. Vancouver, Canada: NeurIPS.
- Nosovskiy, G.V., Liu, D., and Sourina, O., 2008. Automatic clustering and boundary detection algorithm based on adaptive influence function. *Pattern Recognition*, 41 (9), 2757–2776.
- Novikov, A.V., 2019. Pyclustering: data mining library. *Journal of Open Source Software*, 4 (36), 1230.
- Ogbuabor, G., and Ugwoke, F., 2018. Clustering algorithm for a healthcare dataset using silhouette score value. *International Journal of Computer Science and Information Technology*, 10 (2), 27–37.
- Pei, T., et al., 2015. Density-based clustering for data containing two types of points. *International Journal of Geographical Information Science*, 29 (2), 175–193.
- Pei, T., et al., 2006. A new approach to the nearest-neighbour method to discover cluster features in overlaid spatial point processes. *International Journal of Geographical Information Science*, 20 (2), 153–168.
- Perruchet, C., 1983. Constrained agglomerative hierarchical classification. *Pattern Recognition*, 16 (2), 213–217.
- Rao, J., et al., 2020. Q. LSTM-TrajGAN: a deep learning approach to trajectory privacy protection. In: *11th International Conference on Geographical Information Science (GIScience 2021)-Part I*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Vol. 117, 12:1–12:17.
- Rasmussen, C.E., 1999. The infinite Gaussian mixture model. *NIPS*, 12, 554–560.
- Rey, S.J., and Anselin, L., 2010. Pysal: a python library of spatial analytical methods. In: *Handbook of applied spatial analysis*. Berlin: Springer, 175–193.

- Rue, H., and Held, L. 2005. *Gaussian Markov random fields: theory and applications (monographs on statistics and applied probability)*. Boca Raton, FL: Chapman Hall/CRC.
- Scheinberg, K., Ma, S., and Goldfarb, D., 2010. Sparse inverse covariance selection via alternating linearization methods. In: *Advances in neural information processing systems*, 23. Vancouver, Canada: NIPS.
- Soliman, A., et al., 2017. Social sensing of urban land use based on analysis of twitter users' mobility patterns. *PLoS One*, 12 (7), e0181657.
- Steinley, D., 2004. Properties of the Hubert-Arabie adjusted Rand index. *Psychological Methods*, 9 (3), 386–396.
- Sugar, C.A., and James, G.M., 2003. Finding the number of clusters in a dataset: an information-theoretic approach. *Journal of the American Statistical Association*, 98 (463), 750–763.
- Syakur, M., et al., 2018. Integration k-means clustering method and elbow method for identification of the best customer profile cluster. In: *IOP conference series: materials science and engineering*, Vol. 336. Surabaya, Indonesia: IOP Publishing, 012017.
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58 (1), 267–288.
- Tobler, W.R., 1970. A computer movie simulating urban growth in the Detroit region. *Economic Geography*, 46 (sup1), 234–240.
- Viterbi, A.J., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13 (2), 260–269.
- Wang, F., 2020. Why public health needs GIS: a methodological overview. *Annals of GIS*, 26 (1), 1–12.
- Webster, R., and Burrough, P., 1972. Computer-based soil mapping of small areas from sample data: I. multivariate classification and ordination. *Journal of Soil Science*, 23 (2), 210–221.
- Wilson, J.P., et al., 2020. A five-star guide for achieving replicability and reproducibility when working with GIS software and algorithms. *Annals of the American Association of Geographers*, 111 (5), 1311–1317.
- Wolf, L.J., 2021. Spatially–encouraged spectral clustering: a technique for blending map typologies and regionalization. *International Journal of Geographical Information Science*, 35 (11), 2356–2373.
- Wu, L., et al., 2014. Intra-urban human mobility and activity transition: evidence from social media check-in data. *PLoS One*, 9 (5), e97010.
- Xing, H., and Meng, Y., 2018. Integrating landscape metrics and socioeconomic features for urban functional region classification. *Computers, Environment and Urban Systems*, 72, 134–145.
- Yan, B., et al., 2017. From ITDL to Place2Vec: reasoning about place type similarity and relatedness by learning embeddings from augmented spatial contexts. In: *Proceedings of the 25th ACM SIGSPATIAL international conference on advances in geographic information systems*, 1–10. New York, NY: ACM.
- Yang, D., et al., 2014. Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45 (1), 129–142.
- Yuan, N.J., et al., 2015. Discovering urban functional zones using latent activity trajectories. *IEEE Transactions on Knowledge and Data Engineering*, 27 (3), 712–725.
- Zhang, T., Ramakrishnan, R., and Livny, M., 1996. Birch: an efficient data clustering method for very large databases. *ACM SIGMOD Record*, 25 (2), 103–114.
- Zhao, Z., et al., 2016. Understanding the bias of call detail records in human mobility research. *International Journal of Geographical Information Science*, 30 (9), 1738–1762.
- Zhu, A.-X., et al., 2018. Spatial prediction based on third law of geography. *Annals of GIS*, 24 (4), 225–240.
- Zhu, R., Janowicz, K., and Mai, G., 2019. Making direction a first-class citizen of tobler's first law of geography. *Transactions in GIS*, 23 (3), 398–416.
- Zou, H., 2006. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101 (476), 1418–1429.