

Reducing False Discoveries in Statistically-Significant Regional-Colocation Mining: A Summary of Results

Subhankar Ghosh ✉

Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN, USA

Jayant Gupta ✉

Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN, USA

Arun Sharma ✉

Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN, USA

Shuai An ✉

Department of Economics, University of Minnesota, Minneapolis, MN, USA

Shashi Shekhar ✉

Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN, USA

Abstract

Given a set S of spatial feature types, its feature instances, a study area, and a neighbor relationship, the goal is to find pairs \langle a region (r_g), a subset C of S \rangle such that C is a statistically significant regional-colocation pattern in r_g . This problem is important for applications in various domains including ecology, economics, and sociology. The problem is computationally challenging due to the exponential number of regional colocation patterns and candidate regions. Previously, we proposed a miner [8] that finds statistically significant regional colocation patterns. However, the numerous simultaneous statistical inferences raise the risk of false discoveries (also known as the multiple comparisons problem) and carry a high computational cost. We propose a novel algorithm, namely, multiple comparisons regional colocation miner (MultComp-RCM) which uses a Bonferroni correction. Theoretical analysis, experimental evaluation, and case study results show that the proposed method reduces both the false discovery rate and computational cost.

2012 ACM Subject Classification Information systems \rightarrow Data mining; Computing methodologies \rightarrow Spatial and physical reasoning

Keywords and phrases Colocation pattern, Participation index, Multiple comparisons problem, Spatial heterogeneity, Statistical significance

Digital Object Identifier 10.4230/LIPIcs.GIScience.2023.3

Funding This material is based upon work supported by the National Science Foundation under Grants No. 2118285, 2040459, 1901099, and 1916518.

Acknowledgements We also thank Kim Koffolt, Yash Travadi, and the Spatial Computing Research Group for valuable comments and refinements.

1 Introduction

Regional-colocation patterns are (study sub-area R , feature-type subset C) pairs such that instances of feature-types in C often are present in R in close proximity. Given a set S of spatial features (e.g., coffee shops, restaurants), their feature instances, a study area, and a neighbor relationship (e.g., geographic proximity), the goal is to identify pairs \langle region r_g , subset C of S \rangle such that instances of C are statistically significant in that region r_g . Figure 1(a) shows a set of instances input into a regional-colocation miner, consisting of three different spatial feature types, a neighborhood relation between feature instances, and a



© Subhankar Ghosh, Jayant Gupta, Arun Sharma, Shuai An, and Shashi Shekhar; licensed under Creative Commons License CC-BY 4.0

12th International Conference on Geographic Information Science (GIScience 2023).

Editors: Roger Beecham, Jed A. Long, Dianna Smith, Qunshan Zhao, and Sarah Wise; Article No. 3; pp. 3:1–3:18

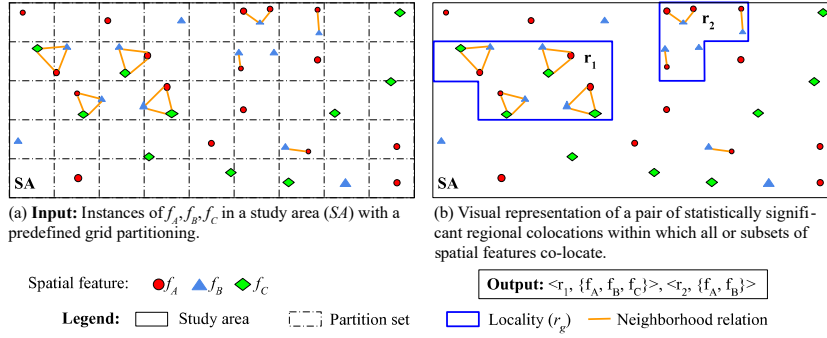
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

3:2 Reducing False Discoveries in Statistically-Significant Regional-Colocation Mining

space partitioning. Figure 1 (b), shows the set of statistically significant regional colocations identified after significance testing (described in Section 2.2). The output is a pair of regional colocations: r_1 showing a strong regional colocation between all three features (i.e., f_A , f_B , and f_C) and r_2 showing a strong regional-colocation between two features (i.e., f_A and f_B). The rest of the area within the map shows less spatial interaction (low participation index) between these features.




■ **Figure 1** Regions where all or subsets of f_A , f_B and f_C significantly co-locate in the study area.

The problem of mining statistically significant regional-colocation patterns is societally important with applications in retail, public health, ecology, public security, transportation, etc. For example, retail establishments (e.g., fast food chains and coffee shops) often colocate to reach each other’s customers. Thus, finding statistically significant regional colocation patterns among competing retail stores has tremendous value for retail analysis. When identifying colocation patterns in societal domains, it’s important to minimize the chance of false discoveries. A famous historical example was between 1900 and 1904 when urban districts of San Francisco experienced an outbreak of bubonic plague, resulting in 119 deaths. The federal and state authorities falsely identified the victims’ ethnicity as a highly correlated feature to the plague. This false discovery brought an immense adverse impact on San Francisco’s management of the plague. Even when we don’t unfairly stigmatize groups or regions, false discoveries waste money, and resources. Comparing the city’s response to the same plague between 1907 and 1908, where rats were correctly identified as a highly correlated feature and the plague was swiftly contained, the negative impact of false discovery was even more strongly felt [13]. Table 1 provides application domains and use cases.

■ **Table 1** Regional-colocation applications.

Application Do-main	Example
Retail	$\langle \text{China}, \{\text{McDonald's and KFC}\} \rangle, \langle \text{USA}, \{\text{McDonald's and Jimmy John's}\} \rangle$
Public Health	$\langle \text{Ports}, \{\text{Plague and rats}\} \rangle, \langle \text{Middle East}, \{\text{Middle East Respiratory Syndrome (MERS) in 2012 and MERS-CoV}\} \rangle$
Ecology	$\langle \text{Indian/Pacific Ocean}, \{\text{Anemone and Clownfish}\} \rangle, \langle \text{Nile River delta}, \{\text{Nile Crocodile and Egyptian Plover}\} \rangle$
Public Safety	$\langle \text{Region around bars}, \{\text{Assault crimes and drunk driving}\} \rangle$
Transportation	$\langle \text{Near bus depots}, \{\text{High } NO_x \text{ concentrations and buses}\} \rangle$



■ **Figure 2** Comparison with Related Work.

The problem of statistically significant regional-colocation pattern detection (*SSRCPD*) is computationally challenging due to the following reasons: (1) Significance testing in this problem requires considering multiple statistical inferences simultaneously which leads to an increase in Type-I error (i.e false discoveries). (2) There is an exponential number of candidate regional patterns, e.g., the dataset used in the case study (Section 6) consists of 1473 different retail brands and their locations in Minnesota, resulting in 2^{1473} different candidate patterns. (3) Spatial partitioning approach would lead to an infinite number of candidate region subsets.

Figure 2 shows a decision tree that distinguishes our manuscript from previous works, where *SSRCPD* refers to *Statistically significant regional colocation pattern detection*. Earlier work on regional-colocation pattern detection either uses data unaware space partitioning (e.g., Quadtree [4, 12]) or clustering of colocation instances [5, 7]. However, these techniques lack statistical significance testing and depend on input parameters (e.g., participation index threshold) which may vary geographically. Statistically significant global colocation mining was introduced by [1], while statistically significant regional colocation mining was first explored in [8]. In [8] we proposed *SSRCM* which utilizes a subgraph enumeration approach to detect statistically significant regional colocation patterns where the regions would be composed of one or more contiguous atomic partitions (smallest region within which a candidate pattern is statistically significant). This algorithm was expensive because expanding the region within which the pattern was statistically significant required recalculating the *p-value*. Since detecting statistically significant regional colocation patterns requires performing multiple simultaneous statistical inferences, this results in the multiple comparisons problem [14], which risks false discoveries (a.k.a. Type-I errors). The problem results in a rapid increase in the probability of Type-I error as the number of partitions increases. To address the multiple comparisons problem, we propose a robust statistically significant regional colocation miner (MultComp-RCM) using a Bonferroni correction [3]. The proposed approach recommends stricter *p-values* to reduce false discoveries (Type-I errors), thus setting an upper bound on the overall significance level (α , which is 0.05 for a 95% statistical confidence).

Contributions.

- We proposed a new approach Multiple comparisons regional colocation miner (MultComp-RCM) to reduce false positives using a well-established statistical technique for multiple comparisons correction, the Bonferroni test.
- The paper provides a comparative analysis showing that the proposed MultComp-RCM is computationally more efficient than SSRCM.
- The paper describes a sensitivity analysis using synthetic data which shows that MultComp-RCM requires an increasingly smaller number of significance tests and participation index computations for an increasing number of regions.
- We proposed a case study on retail establishments in Minnesota using the Safegraph POI dataset [8]. The proposed method discovers new regional-colocation patterns involving fast food and coffee retailer feature-type subsets in a Minnesota counties study area. We also confirm that the Bonferroni correction in our method reduces false discoveries.

Scope. For simplicity, this paper focuses on regional-colocation patterns consisting of two or three different features. In our case study, we enumerated regions based on a contiguous collection of counties. Nevertheless, this work can be extended to different types of regions (e.g., ports). We also do not consider segregation patterns (negative spatial interaction) or the temporal aspects of the patterns.

Organization. The paper is organized as follows. Section 2 reviews basic concepts and formally defines the problem. In section 3 we briefly review SSRM and describe the proposed approach (MultComp-RCM). Section 4 gives a theoretical analysis of MultComp-RCM. We present the experimental evaluation in Section 5 and a case study in Section 6. Section 7 briefly surveys related work and discussion. Section 8 concludes the paper with future work.

2 Basic Concepts and Problem Definition

First, we review basic concepts related to colocation detection, statistical significance testing, and the multiple comparisons problem. Then, we formally define statistically significant regional colocation pattern detection.

2.1 Colocation detection

In this section, we briefly introduce some taxonomy and the basic concept used to define colocation pattern detection with examples. The basic concepts are as follows:

A **feature instance** is a geo-located spatial entity which is a type of Boolean feature f with a geo-reference point location p (e.g., latitude, longitude), represented as $\langle f, p \rangle$. Multiple instances of a feature are represented as f_i and can be related to other feature instances f_j via a **neighbor relation** \mathcal{R} . For example, geographic proximity is represented as $\mathcal{R}_{f_i, f_j} \leq \theta$, where θ is the neighbor relation threshold. In a **neighbor graph**, we represent features that satisfy such relations as a *node* and their relationship as an *edge*.

A **colocation candidate** C is a set of features defined in the given study area (SA) or a sub-region (r_g) where $r_g \in SA$. For example, Figure 1(a) shows 17 spatial objects of type f_A (circle), 12 spatial objects of type f_B (triangle), and 9 instances of colocation pattern $\{f_A, f_B\}$. An instance of a **colocation** satisfies the neighborhood relation \mathcal{R} and forms a **clique**.

A **participation ratio** (pr) is the ratio of feature instances participating in a relation \mathcal{R} to the total number of instances inside the study region (SA). For a given colocation candidate C and feature f , it is represented as $pr(f, C)$ as shown in Equation 1:

$$pr(f, C) = \frac{\text{participating_instances}(f, C)}{\text{instance}(f)}. \quad (1)$$

For the feature instances shown in Figure 1(a) the participation ratio values for the relation $\{f_A, f_B\}$ are $pr(f_A, \{f_A, f_B\}) = \frac{9}{17}$ and $pr(f_B, \{f_A, f_B\}) = \frac{8}{12}$. Further, the participation ratio within a region (r_g) for a feature f is defined as $pr(f, [r_g, C])$. For example, in Figure 1 $pr(f_A, [r_2, \{f_A, f_B\}])$ and $pr(f_B, [r_2, \{f_A, f_B\}])$ in region r_2 and has the value $\frac{4}{4}$ and $\frac{3}{4}$ respectively.

A **participation index** (pi) is the minimal participation ratio of all feature types in a colocation candidate as described in Equation 2:

$$pi(C) = \min_{f \in C} (pr(f, C)). \quad (2)$$

The participation index quantifies the spatial interaction within features. Figure 1(a) shows participation index of features f_A, f_B which can be represented as $pi(\{f_A, f_B\})$ which is $min(\frac{9}{17}, \frac{8}{12})$ or $\frac{8}{17}$. A **regional participation index** is the minimal participation ratio of all feature types in the colocation candidate C within region r_g as shown below

$$pi([r_g, C]) = \min_{f \in C}(pr(f, [r_g, C])) \quad (3)$$

For instance in Figure 1, $pi([r_2, \{f_A, f_B\}]) = min(\frac{4}{4}, \frac{3}{4}) = \frac{3}{4}$.

Colocation patterns [16] is the set of prevalent colocation candidates (based on a prevalence measure, e.g. pi), i.e., candidates comprised of features having a high positive spatial interaction. A **regional-colocation pattern** [12] is a paired region (r_g) and colocation pattern (C), i.e., $\langle r_g, C \rangle$ where the features in pattern C have a high positive spatial interaction in r_g .

2.2 Statistical Significance in Colocation Detection

A statistically significant colocation determines whether an assigned positive spatial interaction between features is statistically significant or could have been observed if the features were in complete spatial randomness (CSR). Other properties in CSR are as follows:

- Every feature instance has an equal probability of existing at any point in the study area.
- The locations of any feature instances in the study area are independent of each other.

A **null hypothesis** (H_0) is a statement of “no effect” or “no difference”. In our problem, the null hypothesis represents the scenario under which there is no spatial interaction between the features in the dataset, i.e., their existence is completely independent of each other.

An **alternative hypothesis** (H_a) is a statement that is tested against a null hypothesis. In our problem, an alternative hypothesis represents the scenario under which there is a positive spatial interaction between the features in the dataset in a region of interest.

A **Type-I error** refers to the erroneous rejection of an actually true null hypothesis (or a false positive). In our problem, this would refer to incorrectly assigning a candidate regional-colocation pattern as statistically significant, even though there is a high probability of this pattern being found in CSR or H_0 .

A **Type-II error** refers to the failure to reject a null hypothesis (H_0) that is actually false (or a false negative). This would translate into incorrectly assigning a candidate regional-colocation pattern as not statistically significant.

A **point distribution** is a collection of geo-distributed points referring to an event (e.g., road accident) in a spatial domain. A **point process** (PP) is a statistical process that defines the probability distribution of a point over a region. Point processes are essential for defining the null or alternative hypothesis for our statistical significance test.

A **Poisson point process** is defined in a generalized space S_P with intensity Λ having the following properties:

1. The number of points in a bounded Borel set (bounded sets that can be constructed from open or closed sets by repeatedly taking countable unions and intersections) $B \subset S_P$ is a Poisson random variable with mean $\Lambda(B)$.
2. The number of points in n disjoint Borel sets forms n independent random variables. This property results in independent scattering or complete independence.

Null hypothesis generation.

- For an identical distribution, we generate an equal number of instances of each feature in every partition using summary statistics of the constituent features of the pattern. This ensures that the null hypotheses datasets (although in CSR) closely model the observed dataset in each atomic partition.

- For independence, we sample instances from a Poisson point process [11]. To check for acceptable auto-correlation, we use a pair correlation function (PCF) or $g(d)$ up to a distance d , where d is data-driven. When $g(d) > 1$, it suggests there is clustering at a distance d within the feature instances, while $g(d) = 1$ represents CSR.

Statistical significance test. Since the participation index (pi) is used to quantify the strength of spatial interaction, the objective is to determine the probability of a pattern's pi in the observed data. Let $pi_{i\emptyset}(C)$ denote the participation index for pattern C in the null hypothesis and $pi_{obs}(C)$ represent the participation index for candidate colocation C in the observed data. Then, we compute the following probability [1]:

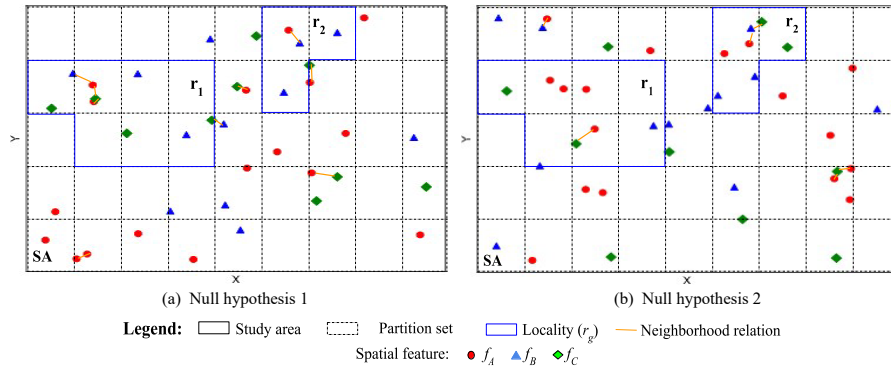
$$p = pr(pi_{i\emptyset}(C) \geq pi_{obs}(C)) = \frac{R^{\geq pi_{obs}} + 1}{R + 1} \tag{4}$$

where $R^{\geq pi_{obs}}$ represents the number of Monte Carlo simulations within which the participation index ($pi_{i\emptyset}(C)$) for pattern C is greater than in the observed data ($pi_{obs}(C)$) and R refers to the total number of Monte Carlo simulations. If $p \leq \alpha$, we consider $pi_{obs}(C)$ as statistically significant at level α .

Regional statistical significance test. To test for regional significance, we use simulated (i.e., computer-generated) candidate regions. For example, if we are trying to determine if f_A and f_B are statistically significant in locality r_2 (Figure 1(b)), we generate null hypothesis samples within the locality's boundaries and use the participation index result from each sample to perform the significance test for r_g . Figure 3(a) and 3(b) display two of the R different null hypotheses which are used to compare the participation index of the regional-colocation pattern $[r_2, \{f_A, f_B\}]$ in the observed data. As shown, the participation index for the pattern $[r_2, \{f_A, f_B\}]$ is $1/3$ in both null hypotheses respectively. For a statistical confidence level of 95% the following inequality should hold:

$$\left[\sum_{r=1}^{R=99} \mathbb{1}(pi_{obs}([r_g, \{f_A, f_B\}]) \leq pi_{i\emptyset_r}([r_g, \{f_A, f_B\}])) \right] < 5 \tag{5}$$

where r_g is the region of interest, and $\mathbb{1}$ is an indicator function. We can compute R from $\alpha = 0.05$ using $\alpha(R + 1) = 5$ [2].



■ **Figure 3** Two example null hypotheses for significance testing of the observed data in Figure 1.

The **multiple comparisons problem** [14] occurs when every inference in a set of statistical inferences simultaneously has the potential to produce a discovery. The more inferences are made on a particular data set, the more likely it is to incorrectly reject the null hypothesis.

Most techniques to address this problem require a stricter significance threshold for individual comparisons to compensate for the number of inferences being made. A stated confidence level generally applies to individual tests. It is often desirable to have a confidence level for a whole family of simultaneous tests.

The **Bonferroni correction** [3] is a method to address the multiple comparisons problem and the simplest method for reducing Type-I errors. It is a conservative method with a greater risk of failure to reject a false null hypothesis, thus resulting in Type-II errors.

2.3 Formal problem formulation

The problem of statistically significant regional-colocation pattern detection is as follows:

Input:

1. A set (F) of spatial-features
2. N geo-located spatial feature instances.
3. A study area S_A composed of space partitions (e.g., counties).
4. A statistical significance level α .
5. A neighbor relationship (\mathcal{R}).

Output: Statistically significant regional-colocation patterns, $\langle r_g, C \rangle$ where $C \subset F$.

Objective: Reducing Type-I error (false positives).

Constraints: Higher statistical confidence of output patterns.

Reasoning behind problem output. Testing for statistical significance on regional-colocation outputs ensures that spurious patterns aren't detected from the dataset. Otherwise, regions may be enumerated due to a high density of feature instances or spatial auto-correlation. In addition, significance testing for the union of many partitions leads to multiple statistical inferences. Due to the union of partitions, the probability of finding chance patterns within the bigger region (i.e., the union of partitions) is higher; this phenomenon is not accounted for by the p -value threshold for a single partition. This leads to the multiple comparisons problem, resulting in a higher false discovery rate. In application domains related to regional-colocation pattern detection, reducing Type-I errors (false positives) takes higher priority over reducing Type-II errors (false negatives). These Type-II errors might result in missing the detection of certain patterns which might have a lower p -value.

In this situation, checking for a particular α level in the individual statistical inferences is insufficient. We also need to control the family-wise error rate which represents the probability of making one or more false discoveries (Type-I errors) [14]. We use the Bonferroni correction in MultComp-RCM to tackle this problem arising from multiple hypothesis tests. This conservative approach ensures that the pattern output has high statistical confidence while ignoring patterns that might have comparatively lower confidence, which is our primary objective. Another benefit of this method is the computational efficiency due to the smaller number of significance tests and participation index computations required as compared to the baseline [8]. The Bonferroni correction proposes stricter p -value thresholds which might be a bottleneck for large scale applications, such as when dealing with hundreds of atomic partitions. This may also lead to a higher possibility of false negatives (Type-II errors).

3 Methodology

To keep the paper self-contained, we first briefly review the SSRCM, our previous statistically significant regional-colocation miner [8], and a sub-routine on significance testing. We then describe the proposed approach in Section 3.2 and provide an example highlighting the computational cost savings of the new approach.

3.1 Statistically Significant Regional-Colocation Miner

Key idea. In [8], we started by considering partitions with at least 3 instances of each feature which comprise the regional-colocation pattern. This ensures that the features constituting the pattern all have a considerable presence in the enumerated partitions. We then use the regional statistical significance test as described in Algorithm 1 to determine the atomic footprints of the pattern, i.e., statistically significant pattern within individual partitions. While computing the participation index, we limit our neighborhood to an empirically determined distance (d) to mine meaningful collocated features.

■ **Algorithm 1** Significance testing.

Input:

- A spatial dataset S consisting of features $\{f_A, f_B, \dots\}$
- A study area (S_A) and an atomic partition $r_g \subset S_A$
- Statistical significance level α
- A candidate colocation pattern C
- A set of R Null hypotheses (NH_\emptyset) data each modelled as colocation C in atomic partition r_g
- Distance d for participation index (pi) calculation

Output:

1. $\langle r_g, C \rangle$ is significant or not
2. $p\text{-value}_C^{r_g}$

1: **procedure** SIGNIFICANCE TESTING
2: Statistically significant result $SSR_C^{r_g} \leftarrow \text{False}$
3: Counter $R^{\geq pi_{obs}} \leftarrow 0$
4: Calculate pi_{obs} for C at d in r_g
5: **for** $i \in [1, R]$ **do**
6: Calculate the $pi_{\emptyset, i}$ of C at d in the i^{th} NH_\emptyset
7: **if** $pi_{\emptyset, i} \geq pi_{obs}$ **then**
8: $R^{\geq pi_{obs}} \leftarrow R^{\geq pi_{obs}} + 1$
9: $p\text{-value}_C^{r_g} = \frac{R^{\geq pi_{obs}} + 1}{R + 1}$
10: **if** $p\text{-value}_C^{r_g} \leq \alpha$ **then**
11: $SSR_C^{r_g} \leftarrow \text{True}$ ▷ (i.e., $\langle r_g, C \rangle$ is statistically significant)
12: **else**
13: $SSR_C^{r_g} \leftarrow \text{False}$ ▷ (i.e., $\langle r_g, C \rangle$ is not statistically significant)
14: **return** $SSR_C^{r_g}, p\text{-value}_C^{r_g}$

For finding the union of partitions, we first form an undirected unweighted graph ($G = (V, E)$) where each vertex (V) refers to a partition within which the pattern is statistically significant, and an edge (E) between two V s represents a shared boundary between them. The graph representation allows the use of graph traversal algorithms (e.g., DFS) to find statistically significant regions which are the union of partitions in V .

We note that the union of two atomic footprints within which a candidate regional colocation pattern is statistically significant does not imply that the resultant footprint is a significant regional-colocation pattern. Thus, we need to recompute the pi for the candidate pattern in the new region and perform the significance test again. As we progress along the edges of G , the final output is a larger region composed of contiguous atomic partitions such that the candidate pattern is statistically significant, both within the atomic partitions as well as in the region formed by the union of the output atomic footprints. This is represented by the **largest connected component**. Algorithm 2 provides the pseudo-code of $SSRCM$ to find statistically significant regional colocations.

Algorithm 2 Statistically Significant Regional-Colocation Miner (SSCRM).

Input:

- A Spatial dataset S consisting of features $\{f_A, f_B, \dots\}$
- A study area (S_A) and a space partitioning R_g
- Statistical significance level α
- Maximum pattern size N
- Lower bound LB (in meters)
- Upper bound UB (in meters)

Output:

1. List of statistically significant regional colocation patterns [$\langle r_g, C \rangle$]

Variables:

Distance between feature instances d

```

1: procedure STATISTICALLY SIGNIFICANT REGIONAL-COLOLOCATION MINER
2:   for each:  $f_k$  in  $\{f_A, f_B, \dots\}$  do
3:     Generate  $R$  null hypotheses ( $NH_\emptyset$ ) using summary statistics in each  $r_g \in R_g$ .
4:   for each: candidate pattern  $C_m \in \{C_1, C_2, \dots, C_M\}$  do
5:     for distance  $d \in [LB, LB + 10, \dots, UB]$  do
6:       for each:  $r_g \in R_g$  do
7:          $SSR_{C_m}^{r_g}, p\text{-value} \leftarrow$  Significance Testing( $S, r_g, \alpha, C_m, NH_\emptyset, d$ )
8:         if  $SSR_{C_m}^{r_g}$  is True then
9:           Insert  $r_g$  in significant atomic partitions list
10:        Compose Neighborhood graph ( $G$ ) from significant atomic partitions list
11:         $r_g^{final} \leftarrow r_g^{maxPI}$   $\triangleright$  atomic partition in  $G$  with highest  $pi$ 
12:        for each:  $r_g \in$  Depth First Graph Traversal of  $G$ 
13:          from vertices adjacent to  $r_g^{maxPI}$  do  $\triangleright r_g \neq r_g^{maxPI}$ 
14:             $r_g^{temp} \leftarrow r_g^{final} \cup r_g$ 
15:             $SSR_{C_m}^{r_g}, p\text{-value} \leftarrow$  Significance Testing( $S, r_g^{temp}, \alpha, C_m, NH_\emptyset, d$ )
16:            if  $SSR_{C_m}^{r_g}$  is True then
17:               $r_g^{final} \leftarrow r_g^{temp}$ 
18:            Add  $\langle r_g^{final}, C_m \rangle$  to [ $\langle r_g, C \rangle$ ]
19:   return [ $\langle r_g, C \rangle$ ]
  
```

3.2 Multiple Comparisons Regional Colocation Miner (MultComp-RCM)

Key Idea. We observe that the baseline *SSRCM* computes a significance test for every union of statistically significant partitions, resulting in many participation index (pi) computations and significance tests. We address this by using a Bonferroni correction, which selects atomic partitions conservatively, increasing the chances that their union is also statistically significant. The Bonferroni correction reduces the need to perform a regional statistical significance test for each union operation.

A Bonferroni correction is used when several independent statistical inferences are being performed simultaneously. Although a given significance threshold (α) on the p -value may be appropriate for an individual test, it is not sufficient for the set of all comparisons. To reduce many false positives, the α needs to be lowered to account for the number of comparisons performed. The Bonferroni correction sets the statistical significance threshold for the entire set of n comparisons to α/n or, equivalently, by multiplying the p -value by n , and then applying the standard threshold α . This conservative correction works even under the most extreme circumstances (e.g., when all n tests are independent of one another).

In the proposed approach we check for statistical significance in each input partition. Then, we perform a graph traversal starting from the atomic partition with the highest pi value for the candidate regional-colocation pattern. Then, instead of recomputing the pi and testing the candidate pattern in the new bigger region (composed of atomic partitions) for

3:10 Reducing False Discoveries in Statistically-Significant Regional-Colocation Mining

statistical significance, we perform a Bonferroni correction. Thus if we were initially checking for a threshold level of 0.05, then for the union of two partitions, we would be checking for a threshold level of $0.05/2$ in each atomic partition. This conservative threshold reduces Type-I error by returning regions with much higher statistical confidence. The union of the atomic partitions is sequential and every atomic partition must satisfy the adjusted p-value threshold to be considered for the union.

Algorithm 3 provides a snippet of MultComp-RCM showing the use of the Bonferroni correction. Lines 13-18 show the new steps in the refined approach.

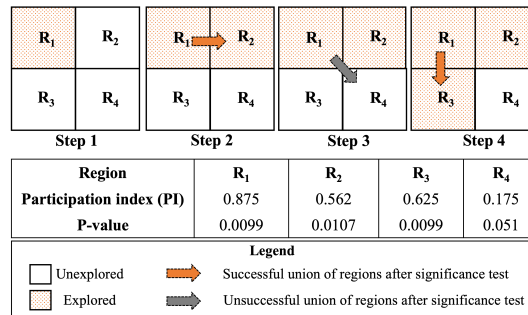
■ **Algorithm 3** MultComp-RCM snippet.

```

1: procedure MULTCOMP-RCM
2:   ⋮
12:   $n \leftarrow 1$  ▷ Number of atomic partitions in the region.
13:  for each:  $r_g \in$  Depth First Graph Traversal of  $G$ 
14:    from vertices adjacent to  $r_g^{maxPI}$  do ▷  $r_g \neq r_g^{maxPI}$ 
15:    Subgraph ( $SG$ )  $\leftarrow r_g^{final} \cup r_g$ 
16:    flag  $\leftarrow 1$ 
17:    flag = BONF_CHECK(flag,  $SG$ ,  $n$ )
18:    if flag == 1 then
19:      Update  $r_g^{final} \leftarrow SG$ 
20:       $n \leftarrow n + 1$ 
21:    Add  $\langle r_g^{final}, C_m \rangle$  to  $\langle r_g^C, C \rangle$ 
22:    ⋮
23: procedure BONF_CHECK(flag,  $SG$ ,  $n$ )
24:   $p\text{-value}_{threshold} \leftarrow \alpha / (n + 1)$ 
25:  for each:  $node \in SG$  do
26:    if  $p$ -value of pattern  $C_m$  in  $node \not\leq p\text{-value}_{threshold}$  then
27:      flag  $\leftarrow 0$ 
28:  return flag

```

Figure 4 shows an execution trace of merging 4 neighboring partitions. Each region has a participation index and a p -value computed individually. Then, Steps 1-4 show the process of combining these partitions based on either additional statistical significance tests and participation index computations (for SSRM) or using a tighter p -value threshold for MultComp-RCM. Table 2 compares the number of computations for the two approaches and clearly shows the lower computational requirements of MultComp-RCM. When performing the union of two regions using MultComp-RCM, the new threshold as per the Bonferroni correction is applied to each of the two regions (as in procedure BONF_CHECK in Algorithm 3) for a successful union.



■ **Figure 4** Execution trace of SSRM and MultComp-RCM.

■ **Table 2** Comparing the cumulative number of statistical significance tests ($C\#$), participation index computation (pi cal.), and p-value thresholds (p -val th.) between $SSRCM$ (denoted as S) and $MultComp$ -RCM (denoted as R).

Steps	$C\#S$	$C\#R$	pi cal. S	pi cal. R	p -val. th. S	p -val. th. R
0	4	4	4	4	0.05	0.05
1	5	4	5	4	0.05	0.05
2	6	4	6	4	0.05	0.025
3	7	4	7	4	0.05	0.0167
4	8	4	8	4	0.05	0.0125

4 Theoretical Analysis

► **Lemma 1.** *MultComp – RCM has lower or equal Type-I error than SSRCM.*

Proof. Algorithm 1 called by Algorithm 2 and 3 in line 7 extracts atomic partitions within which a regional-colocation pattern is statistically significant.

The Bonferroni correction in procedure BONF_CHECK in Algorithm 3 controls the experiment-wide false positive rate (π) by specifying the significance level (α) for each test, where a test is significant if p -value $\leq \alpha$. The probability of no Type I error (false positives) in n independent tests is $(1 - \alpha)^n$, if each test is at level α . Therefore, the probability of at least one false positive π is $1 - (1 - \alpha)^n$. For an experiment-wide false positive rate of π , the α for each test should be $\alpha = 1 - (1 - \pi)^{1/n}$. Using binomial approximation, $(1 - \alpha)^n \simeq 1 - n\alpha$, which gives $\alpha = \pi/n$. For an experiment-wide false positive value $\pi = 0.05$, the α (false positive rate for each test) should be less than π , i.e. $\alpha \leq \pi$. Therefore each region and sub-region output by MultComp-RCM has lower Type-I and a precision close to 1. ◀

► **Lemma 2.** *MultComp-RCM has lower or equal computational cost than SSRCM for all observed data, where Bonferroni-revised p-values eliminate lower confidence candidates considered by the original p-value, i.e. $Cost_{MultComp-RCM} \leq Cost_{SSRCM}$.*

Proof. Let $C_{pi}(d)$ be the complexity of participation index (pi) computation for a specific region (dependent on the data d). Let $C_{st}(pi_{obs}, pi_{null}, d)$ be the complexity of significance testing for a specific region (dependent on the pi in observed data d and the null hypothesis). Assume N_1 is the number of space partitions/regions in the dataset, N_2 is the number of space partitions extracted from Algorithm 1, and $N_2 \leq N_1$. Further, assume d_1 is the initial dataset and d_2 is the dataset in each iteration in the SSRCM. Then, the cost of $SSRCM$ is


$$N_1(C_{pi}(d_1) + C_{st}(pi_{obs}, pi_{null}, d_1)) + N_2(C_{pi}(d_2) + C_{st}(pi_{obs}, pi_{null}, d_2)) \quad (6)$$

By contrast, the cost of the proposed MultComp-RCM approach is only $N_1(C_{pi}(d_1) + C_{st}(pi_{obs}, pi_{null}, d_1)) + N_2$. Here, N_2 represents the number of significant partitions for which the p-value needs a comparison against the threshold obtained from Bonferroni correction. ◀

5 Experimental Evaluation


We had three goals for the experiments: (1) To compare the time taken by $SSRCM$ and MultComp-RCM with varying numbers of regional-colocation instances, varying number of atomic partitions, and change in the number of feature instances. (2) To compare the number of significance tests, pi calculations for a varying number of regions. (3) To compare solution quality between $SSRCM$ and MultComp-RCM.

Experiment design. Figure 5 shows the overall validation framework. The metric for comparing the solution quality of *SSRCM* with *MultComp-RCM* was the false positive rate (FPR), while the runtime comparisons were based on the execution time (in seconds) of the individual algorithms. The experiments were done on both real (Safegraph POI) and synthetic data to perform both comparative and sensitivity analysis.



■ **Figure 5** Overall validation framework.

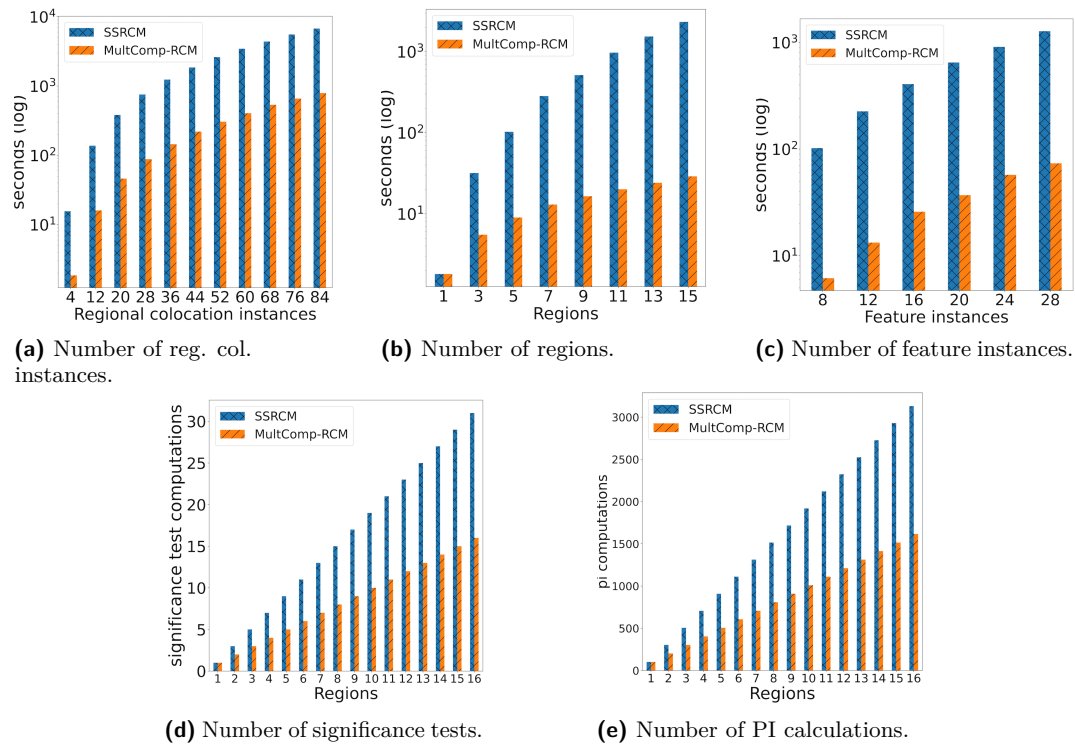
Synthetic data generation. We began with a space partitioning (R_g), a maximum union (or traversal) of regions (L_{max}), and a number of regional-colocation patterns i.e., pairs of $\langle r_g, C \rangle$. We then generated reference points within the partitions using the Poisson point process. At each reference point, we generated circles of diameter d_g which was determined empirically for each region in R_g in the observed dataset. The diameter signifies the smallest distance between features in a colocation C at which they become statistically significant regional colocations. We populated each circle with instances of C . We note that the circles were only used to place collocated instances in a region and were not separate partitioning. Figure 6 shows the process of synthetic data generation.



■ **Figure 6** Synthetic data generation process.

Comparative Analysis. Figure 7a shows the time taken (in log scale) for different regional-colocation instances. For this experiment, we varied the number of regional-colocation instances in each atomic partition from 4 to 84 while keeping other parameters (like the number of regions) constant and record the execution time of both algorithms. Figure 7b compares the execution time with a varying number of atomic partitions (or regions) while keeping the number of regional-colocation instances in each partition constant. Figure 7c shows the time taken with a varying number of feature instances (which constitute the regional-colocation pattern) in each region while keeping the number of regions constant. In all experiments, *MultComp-RCM* is much faster than the baseline *SSRCM*. These results are consistent with Lemma 2, which says that $Cost_{MultComp-RCM} \leq Cost_{SSRCM}$.

Sensitivity Analysis. Figure 7d shows the number of significance tests performed by both algorithms with varying number of regions, while keeping the number of regional-colocation instances constant in each partition. Figure 7e shows the number of participation index



■ **Figure 7** MultComp-RCM outperforms SSRCM [8].

computations performed with varying number of regions with the same constant parameters as above. In both cases, the proposed MultComp-RCM requires lesser number of significance tests and participation index computations for an increasing number of regions.

Solution Quality. We performed controlled experiments on synthetic datasets to compare the solution quality of *MultComp-RCM* with *SSRCM*. Metric for comparison was the false positive rate (FPR).

$FPR = \frac{FP}{FP+TN}$, where FP is the number of false positives, and TN is the number of true negatives. Table 3 shows the experiment results. As shown MultComp-RCM exhibits a lower rate for false pattern discovery than *SSRCM*. This is mainly because MultComp-RCM eliminates regions which barely pass the atomic significance test (borderline statistical confidence) in Algorithm 1, which *SSRCM* fails to reject in the final output.

■ **Table 3** MultComp-RCM generates less false positives.

Pattern	SSRCM False Positive Rate	MultComp-RCM False Positive Rate
A, B, C	0.15	0.03
A, B	0.17	0.01
B, C	0.14	0.01
A, C	0.19	0.04

6 Case Study

We extended our previous case study [8] to show the effectiveness of the proposed approach.

Dataset. We used data from SafeGraph, a mobility data vendor who provides anonymized aggregated location data to researchers studying the effects of COVID-19 on citizen mobility patterns towards numerous Points Of Interest (POIs). The dataset consists of 1473 retail brands in Minnesota. Experiments were performed on colocation patterns consisting of two (e.g., Jimmy John’s, McDonald’s) or three (e.g., Jimmy John’s, McDonald’s, Subway) features. Our null hypothesis generation followed the procedure described in Section 2.2.

Case Study Results. The pattern $C := \{Jimmy\ John's, McDonald's, Subway\}$ was found to be statistically significant when the distance between feature instances was about 1400 meters. The regional footprint was the union of “Dakota” and “Hennepin” Counties. The pi values in the counties were 0.34 and 0.45 respectively. The p -value for the pattern within the counties were 0.02 and 0.01, satisfying the p -value threshold of $\frac{0.05}{2}$ as per the Bonferroni correction for the two partitions. A few additional significant patterns are shown in Table 4 (values rounded to two decimal places).

■ **Table 4** Regional-colocation patterns found to be statistically-significant at distance d .

Colocated features	Counties (<i>participationindex</i> , p – value)	d
{Caribou coffee, Starbucks}	Hennepin (0.34, 0.01)	200 m
{Caribou coffee, Starbucks}	Carver (0.5, 0.02), Hennepin (0.51, 0.01), Washington (0.41, 0.01)	400 m
{Caribou coffee, Starbucks, Dunn Bros}	Hennepin (0.52, 0.01)	1900 m
{Caribou coffee, Starbucks, Dunn Bros}	Hennepin (0.72, 0.01), Washington (0.36, 0.02)	3000 m
{Jimmy John’s, McDonald’s}	Hennepin (0.39, 0.01)	500 m
{Jimmy John’s, McDonald’s}	Dakota (0.36, 0.02), Hennepin (0.51, 0.01)	700 m
{Jimmy John’s, McDonald’s, Subway}	Dakota (0.34, 0.02), Hennepin (0.45, 0.01)	1400 m
{Jimmy John’s, McDonald’s, Subway}	Dakota (0.47, 0.02), Hennepin (0.57, 0.01), Washington (0.43, 0.02)	1500 m

In our previous paper [8], we compared SSRCM with the Quad and QGFR algorithms [12] whose data-aware space partitioning approach is based on the minimum orthogonal bounding rectangle (MOBR). We found that the MOBR-based approach with a participation index threshold of 0.6 produced 3368 potential localities for the pattern $\{r_g, [\text{Caribou Coffee, Starbucks}]\}$. With a confidence level of 95%, MOBR-based approach resulted in 2917 significant and 451 non-significant patterns. Hence, a regional-colocation miner without statistical significance may enumerate output regions where colocations occurred by chance.

7 Related Work and Discussion

Related Work. The concept of colocation was introduced by Shekhar et al. [16]. Huang et al. [10] provided extensive experiments and rigorous discussions regarding the topic and the participation index as a prevalence measure between constituent features. Later, Barua et al. [1] introduced statistical significance testing in global colocation and segregation pattern detection to avoid enumeration of chance patterns in the dataset for both aggregation and segregation patterns but did not mention patterns that are regional (or local). Regional colocation with minimum orthogonal bounding rectangle (MOBR) based approach was studied

by Li et al. [12] while [17] and [4] focused on shapes and zonal patterns, respectively. These methods utilized a threshold on the participation index (pi) without statistical significance testing, leading to the detection of spurious patterns (as discussed in [8]). We [8] recently proposed a subgraph-based approach that incorporates statistical significance in detecting regional colocation patterns. This approach reduced the number of spurious patterns detected by previous methods. However, due to a large number of simultaneous statistical inferences, an increase in false discoveries is also observed. Besides, other patterns [15] and several statistical significance and false discovery reduction techniques have been studied in association rule mining [18, 6]. However, these approaches do not address the inherent variability in spatial data (i.e., different summary statistics of features in each atomic partition). To find subgroups of items, which are generally observed to be statistically significant associations, they compare a quality measure (which assigns to each itemset a numeric value) on the subgroup against that in a statistical model (which corresponds to the null hypothesis). These null hypotheses for significance testing are uniform and do not address spatial variability. Thus these approaches are not directly applicable to regional colocation patterns (more details in Appendix A).

8 Conclusion and Future Work

In this paper, we refined the problem of the statistically significant regional-colocation pattern (*SSCRP*). We proposed a robust *MultComp-RCM* approach that reduces the number of false positives using a Bonferroni correction. We theoretically show that *MultComp-RCM* has a lower or equal Type-I error and computational cost than *SSRCM* along with experimental results. We extended the previous case study on retail establishments in Minnesota using the proposed approach showing a contrast between significant and non-significant patterns.

Future Work. We plan to explore other methods to reduce Type-I errors (false positives) while also addressing Type-II errors (false negatives) arising from the conservative Bonferroni correction approach and further add temporal dimension to these patterns.

References

- 1 Sajib Barua and Jörg Sander. Mining statistically significant co-location and segregation patterns. *IEEE TKDE*, 26(5):1185–1199, 2013.
- 2 Julian Besag and Peter J Diggle. Simple monte carlo tests for spatial pattern. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 26(3):327–333, 1977.
- 3 Carlo Bonferroni. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze*, 8:3–62, 1936.
- 4 M Celik et al. Zonal co-location pattern discovery with dynamic parameters. *ICDM*, 2007.
- 5 Min Deng et al. Multi-level method for discovery of regional co-location patterns. *IJGIS*, 2017.
- 6 Wouter Duivesteijn and Arno Knobbe. Exploiting false discoveries—statistical validation of patterns and quality measures in subgroup discovery. In *2011 IEEE 11th International Conference on Data Mining*, pages 151–160. IEEE, 2011.
- 7 Christoph F. Eick, Rachana Parmar, et al. Finding regional co-location patterns for sets of continuous variables in spatial datasets. In *SIGSPATIAL*, 2008.
- 8 Subhankar et. al. Towards geographically robust statistically significant regional colocation pattern detection. In *Proceedings of the 5th ACM SIGSPATIAL GeoSIM*, pages 11–20, 2022.
- 9 Yan Li et al. Cscd: Towards spatially resolving the heterogeneous landscape of mxif oncology data. In *Proceedings of the 10th ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data*, BigSpatial '22, pages 36–46, New York, NY, USA, 2022. ACM.

- 10 Yan Huang et al. Discovering colocation patterns from spatial data sets: a general approach. *IEEE TKDE*, 16(12):1472–1485, 2004.
- 11 Janine Illian, Antti Penttinen, et al. *Statistical analysis and modelling of spatial point patterns*, volume 70. John Wiley & Sons, 2008.
- 12 Yan Li and Shashi Shekhar. Local co-location pattern detection: a summary of results. In *GIScience*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 13 Guenter B Risse. “A long pull, a strong pull, and all together”: San francisco and bubonic plague, 1907-1908. *Bulletin of the History of Medicine*, 66(2):260–286, 1992.
- 14 G Rupert Jr et al. *Simultaneous statistical inference*. Springer Science & Business Media, 2012.
- 15 Arun Sharma, Jayant Gupta, and Subhankar Ghosh. Towards a tighter bound on possible-rendezvous areas: preliminary results. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pages 1–11, 2022.
- 16 Shashi Shekhar and Yan Huang. Discovering spatial co-location patterns: A summary of results. In *Intl. symposium on spatial and temporal databases*, pages 236–256. Springer, 2001.
- 17 Song Wang et al. Regional co-locations of arbitrary shapes. In *SSTD*, 2013.
- 18 Geoffrey I Webb. Discovering significant patterns. *Machine learning*, 68(1):1–33, 2007.
- 19 David WS Wong. The modifiable areal unit problem (maup). In *WorldMinds: geographical perspectives on 100 problems: commemorating the 100th anniversary of the association of American geographers 1904–2004*, pages 571–575. Springer, 2004.

In this appendix, we address the following questions:

A Why can’t we use existing false discovery reduction techniques from local pattern mining?

Existing techniques for reducing false discoveries in local pattern mining cannot be applied to this problem, because of spatial variability (i.e. constituent features of a regional colocation pattern might have different summary statistics in different atomic partitions). Webb [18] proposed a holdout approach where one divides the data into exploratory and holdout sets. Patterns are generated using the exploratory data, while statistical tests are performed on the generated patterns using the holdout data. This technique may apply to atomic partitions with a large presence of constituent features (e.g., partition T41 in Figure 8). However, it would be counterproductive in partitions where the number of feature instances is very low (e.g., partition T42 in Fig. 8). In such partitions splitting the data points into exploratory and holdout sets would result in very few instances for the pattern detection process.




Figure 8 Feature instances exhibit spatial variability within atomic partitions.

B Why can't this problem be cast as a modified version of frequent itemset mining?

In frequent itemset mining, the task is to find subgroups of items that often occur together in a transaction, e.g., laptop and antivirus software. Previous works have been done on addressing false discoveries in this problem [6]. Such approaches assign the association in the mined subgroup as the alternate hypothesis while the null hypothesis is formulated using a randomized baseline subset. Thus these approaches do not address the independent relationship between hypotheses in different spatial partitions in our problem. As noted earlier, in regional colocation pattern detection, different features might have different summary statistics in different atomic partitions. To model the complete spatial randomness of these features, we generate the null hypotheses in each atomic partition as per the summary statistics of the said features in that specific partition. Thus the null hypothesis generated for the features in one atomic partition is independent of the null hypothesis in other atomic partitions. Therefore, the problem of regional colocation pattern detection cannot be considered a modified version of subgroup discovery in frequent itemset mining.

C How does spatial colocation mining differ from association rule mining?

Data mining techniques have been widely developed to solve challenging problems in various domains. Yet, the underlying assumption of these algorithms does not address the problem of spatial variability. This leads to the detection of spurious patterns in spatial data, also known as the modifiable aerial unit problem (MAUP [19]). Colocation pattern detection resembles association rule mining, but the absence of transactions in colocation mining means techniques in association rule mining cannot be used directly to mine colocation patterns.

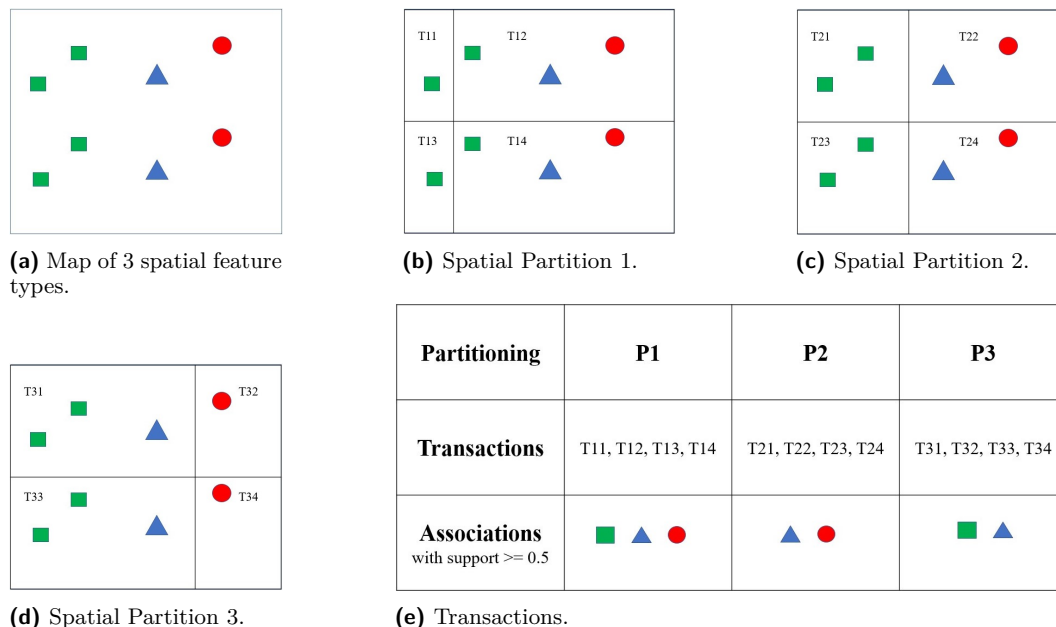







Figure 9 Association rule mining [9] returning different results depending on the spatial partition.

3:18 Reducing False Discoveries in Statistically-Significant Regional-Colocation Mining

Transactions in association rule mining refer to groups of items purchased together. An itemset's support is the fraction of transactions that contain the itemset. Itemsets greater than a user-specified support value yield to the association rule. In spatial data mining, the choice of partition affects the transaction. For example, Figure 9a below shows a dataset with 3 feature types, i.e. <squares>, <triangles>, <circles>. In partition P1 (Figure 9b) <squares, triangles, circles> is a transaction, while in partitions P2 (Figure 9c) and P3 (Figure 9d) <triangles, circles> and <squares, triangles> are the transactions respectively. This is known as the MAUP problem. In colocation pattern detection this is addressed using a neighborhood graph as shown in Figure 10a. A user-defined neighbor relationship R is used to find subsets of features in close geographic proximity. Thus the colocation miner provides a transaction-free approach to mine prevalent patterns.



(a) Neighbor graph based on relation R .

Colocation with Participation Index (PI) ≥ 0.5	PI   = $\min(1/2, 1) = 1/2$	PI   = $\min(1, 1) = 1$
---	---	--

(b) PI of candidate patterns.

■ **Figure 10** Colocation pattern detection [8].