



HAL
open science

Contextual Subgraph Discovery With Mobility Models

Anes Bendimerad, Rémy Cazabet, Marc Plantevit, Céline Robardet

► **To cite this version:**

Anes Bendimerad, Rémy Cazabet, Marc Plantevit, Céline Robardet. Contextual Subgraph Discovery With Mobility Models. COMPLEX NETWORKS 2017, Nov 2017, Lyon, France. pp.477-489. hal-01625068

HAL Id: hal-01625068

<https://hal.science/hal-01625068>

Submitted on 27 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Contextual Subgraph Discovery With Mobility Models

Anes Bendimerad, Rémy Cazabet, Marc Plantevit and Céline Robardet

Abstract Starting from a relational database that gathers information on people mobility – such as origin/destination places, date and time, means of transport – as well as demographic data, we adopt a graph-based representation that results from the aggregation of individual travels. In such a graph, the vertices are places or points of interest (POI) and the edges stand for the trips. Travel information as well as user demographics are labels associated to the edges. We tackle the problem of discovering exceptional contextual subgraphs, i.e., subgraphs related to a context – a restriction on the attribute values – that are unexpected according to a model. Previous work considers a simple model based on the number of trips associated with an edge without taking into account its length or the surrounding demography. In this article, we consider richer models based on statistical physics and demonstrate their ability to capture complex phenomena which were previously ignored.

1 Introduction

The rapid progress of wireless sensor technologies in mobile environments (e.g., GPS, Wi-Fi, RFID) has led to the development of new services for individuals, administrations and companies based on the monitoring of people mobility. For instance, mobility profiles play a central role in context-based search and advertising [3], location modeling [1], traffic planning and route prediction [8], air pollution exposure estimation [5]. The last decade has witnessed a huge growth in the analysis of mobility [7]. These studies focus only on mining trajectories and their applications [13, 14] [16, 19, 20], but do not take into account the contextual information of the individual trajectories. Related contexts of the trajectories are essential data to

Anes Bendimerad and Céline Robardet
Univ Lyon, INSA Lyon, CNRS, LIRIS UMR5205, F-69621 France

Rémy Cazabet and Marc Plantevit
Univ Lyon, Université Lyon 1, CNRS, LIRIS UMR5205, F-69622 France

produce accurate and valuable models of mobility patterns. Recent work has opened this way [2, 12, 9].

Considering a network whose vertices depict places or points of interest (POI) and edges stand for trips and are labeled by sets of transactions¹, that correspond to characteristics of the travels, the problem considered hereafter is to discover and characterize geographical areas that are attractive places and routes for specific contexts. Such areas are frequently visited together in certain contexts, e.g. by users of similar profiles or with trips of similar characteristics. The problem is thus to identify contextual subgraphs (i.e., subgraphs related to a context) that are exceptional in the sens of the Exceptional Model Mining (EMM) approach [11]. EMM aims to build a model on the whole data and a model on data related to a context (generally called pattern), and the exceptionality is assessed by a quality measure that compares the two models. For network analysis, this problem was introduced in [9] where a simple model based on the number of trips associated with an edge was used. In this paper, we study exceptional subgraphs with respect to mobility models that consider into the length or the surrounding demography of the trips. Two mobility models from statistical physics are used: The gravity model [21] and the radiation model [18].

The rest of the paper is organized as follows. Section 2 introduces the problem of exceptional contextual subgraph mining and the related algorithm. The mobility models are defined in Section 3. We report an extensive empirical study in Section 4 and briefly conclude in Section 5.

2 Exceptional contextual subgraphs

We define a graph $G = (V, E, T, \text{EDGE})$ to model urban mobility. V denotes city areas, E denotes the trips from one area to another, and T is a set of transactions of a relation R of schema $S_R = [R_1, \dots, R_p]$. Each attribute R_i takes values in $\mathbf{dom}(R_i)$ that is either nominal or numerical. Thus a transaction $t \in R$ is a tuple (t_1, \dots, t_p) with $t_i \in \mathbf{dom}(R_i)$. EDGE is a mapping of a transaction to an edge: $\text{EDGE} : T \rightarrow E$. For example, if $t_k \in R$ is a transaction such that $\text{EDGE}(t_k) = e_{ij} = (v_i, v_j)$, t_k is a trip from v_i to v_j .

Besides the definition of G , we introduce the notion of *context* and its use to select subgraphs on G . Let a context be a tuple $C = (C_1, \dots, C_p)$ with C_i a restriction on $\mathbf{dom}(R_i)$. C_i can take two different forms depending on the type of R_i :

- If R_i is nominal, $C_i = a$ or $C_i = \star_i$, with $a \in \mathbf{dom}(R_i)$ and $\star_i = \mathbf{dom}(R_i)$
- If R_i is numerical, $C_i = [a, b]$, with $a, b \in \mathbf{dom}(R_i)$, $a < b$.

We say that a transaction (t_1, \dots, t_p) *satisfies or supports* a context C iff $\forall i = 1 \dots p, t_i \in C_i$.

¹ To avoid any ambiguity, notice that the term transaction is different from the one used in Database Management System

Given a context C , the *contextual graph* derived from G is the weighted graph $G_C = (V, E, W_C)$ where $W_C(e)$ is the number of transactions associated to e that satisfy $C = (C_1, \dots, C_p)$:

$$W_C(e) = |\{t = (t_1, \dots, t_p) \mid \text{EDGE}(t) = e \text{ and } \forall i = 1 \dots p, t_i \in C_i\}|$$

Several contexts may be associated to the same contextual graph (i.e. in case of local dependencies between attribute values) and, in such cases, we retain the (unique) most specific one, also called *closed context*. It is defined up to an order relation \preceq defined as follows: A context C^1 is said to be more specific than a context C^2 , denoted $C^1 \preceq C^2$, iff:

- $C_i^2 = \star_i$ or $C_i^1 = C_i^2 = a \in \mathbf{dom}(R_i)$, for R_i a nominal attribute,
- $[a_i^1, b_i^1] \subseteq [a_i^2, b_i^2]$ with $C_i^1 = [a_i^1, b_i^1]$ and $C_i^2 = [a_i^2, b_i^2]$, for all numerical attributes R_i .

A context C is thus *closed* iff $\forall C'$ such that $G_C = G_{C'}$, $C \preceq C'$.

In Exceptional Model Mining (EMM) approach, a pattern (in our case a context) interest is evaluated by its deviation from a null model. An edge $e \in E$ is considered to be exceptional with respect to a context C , if the observed weight $W_C(e)$ is large compared to the expected weight $\widehat{W_C(e)}$. Several discriminative measures can be used. In this paper, we consider the Weighted Relative Accuracy (*WRAcc*) [10] widely used in supervised pattern mining:

$$WRAcc(C, e) = \frac{1}{W_\star(E)} \times (W_C(e) - \widehat{W_C(e)})$$

with $W_\star(E) = \sum_{x \in E} W_\star(x)$ and $\star = (\mathbf{dom}(R_1), \dots, \mathbf{dom}(R_p))$

This measure was also used in [9] in which the expected weight is defined as:

$$\widehat{W_C(e)} = W_\star(e) \times \frac{W_C(E)}{W_\star(E)} \quad (1)$$

This gives the standard definition of the *WRAcc* measure where the expected weight is a portion of the total weight $W_\star(e)$. In the following, this definition of $\widehat{W_C(e)}$ is denoted \mathcal{M}_0 . This mobility model is rather simplistic because it does not take into account the length of trips or the surrounding demographics. In the next section, we propose to use two other mobility models \mathcal{M}_g and \mathcal{M}_r of expected weights.

The mining task we consider is to compute the complete set of closed exceptional graphs, that is to say the couples (C, G_C) where C is closed and G_C is the subgraph made of edges whose associated *WRacc* measure computed with respect to models $\mathcal{M}_0, \mathcal{M}_g$ or \mathcal{M}_r is greater than a threshold *min_threshold* and $W_C(e) > \text{min_weight}$. The algorithm to solve it is presented in Algorithm 1. Cosmic enumerates contexts in a depth-first search manner. In each recursive call, it enumerates an attribute R_i by exploring all its possible values (From Line 6 to Line 28). If all the attributes have been instantiated ($i = p$, Line 2), the subgraph G_C is calculated with respect

to $WRacc$ and \mathcal{M} , if it is exceptional we add it to the result set \mathcal{E} . The detailed description of Cosmic is provided in [9].

Algorithm 1: COSMIc

Input: $C = (C_1, \dots, C_p)$, $G = (V, E, T, \text{EDGE})$, \mathcal{M} the model used to compute the measure $WRacc$ and i the attribute index to be enumerated

Output: \mathcal{E} the set of *exceptional contextual subgraph* patterns under construction

```

1 begin
2   if ( $i = p$ ) then
3     Compute  $G_c$  with respect to  $WRacc$  and  $\mathcal{M}$ 
4      $\mathcal{E} \leftarrow \mathcal{E} \cup (C, G_c)$ 
5   else
6     if ( $R_i$  is symbolic) then
7       for  $a \in \text{dom}(R_i) \cup \{\ast_i\}$  do
8          $C' \leftarrow (C_1, \dots, C_{i-1}, a, C_{i+1}, \dots, C_p)$ 
9          $T'$  is the set of transactions that satisfy  $C'$ 
10         $C''$  is the most specific context that covers all transactions in  $T'$ 
11        if ( $C' = C''$ ) then
12          /*  $C'$  is closed */
13          if  $G_{C'} \neq \emptyset$  then
14            COSMIc( $C', G, i + 1$ )
15        else
16          stack  $\leftarrow ([\min_{a_i \in \text{dom}(R_i)} a_i, \max_{b_i \in \text{dom}(R_i)} b_i], \text{true})$ 
17          while (stack is not empty) do
18            ( $a, b$ ), left  $\leftarrow$  unstack(stack)
19             $C' \leftarrow (C_1, \dots, C_{i-1}, [a, b], C_{i+1}, \dots, C_p)$ 
20             $T'$  is the set of transactions that satisfy  $C'$ 
21             $C''$  is the most specific context that covers all transactions in  $T'$ 
22            if ( $C' = C''$ ) then
23              /*  $C'$  is closed */
24              if  $G_{C'} \neq \emptyset$  then
25                COSMIc( $C', G, i + 1$ )
26              if left = true then
27                interval  $\leftarrow [a, \text{previous}(b)]$ 
28                stack  $\leftarrow$  push(interval, true)
29                interval  $\leftarrow [\text{next}(a), b]$ 
30                stack  $\leftarrow$  push(interval, false)
31        return  $\mathcal{E}$ 

```

3 Mobility models

In mobility modeling [15], urban travels depend on the distances and level of attraction. Thus, the expected amount of transactions between any pair of vertices (v_i, v_j) is not uniform but depends on the distance d_{ij} between v_i and v_j , and on the surrounding populations (or level of attraction) n_i and n_j of these vertices. $W_\star(e_{ij})$ can therefore be considered exceptional, given its associated values of d_{ij} , n_i and n_j , whereas such situations can not be identified by standard $WRacc$ measure. For example, $W_\star(e_{ij})$ can be very large for two points of interest v_i and v_j , while we expect much lower $\widehat{W_\star}(e_{ij})$ with regard to their distance d_{ij} and/or their population n_i and n_j . We propose to model the expected weight of an edge using the gravity or

the radiation models as defined below:

$$\widehat{W_C}(e) = m(e) \times \frac{W_C(E)}{W_*(E)}$$

with $m(e)$ the mobility model, denoted in the following either $g(e)$ for the gravity model, or $r(e)$ for the radiation one.

3.1 The gravity model

In the gravity model, the most widely used mobility model, the number of expected transactions $g(e_{ij})$ between v_i and v_j is defined as [15]:

$$g(e_{ij}) = n_i \times n_j \times f(d_{ij})$$

where $f(d_{ij})$ is known as the *deterrence function* and represents the influence of the distance. In the traditional form of the gravity model, $f(d_{ij})$ is a priori defined as $\frac{1}{d_{ij}^\gamma}$, with γ an optional parameter, usually tuned by regression analysis. However, some recent papers [6, 4] have shown better results using a deterrence function learned from data as follows:

$$f(d) = \frac{\sum_{x,y|d_{xy}=d} W_*(e_{xy})}{\sum_{x,y|d_{xy}=d} n_x \times n_y} \quad (2)$$

In the rest of this article, we will refer to this model as *gravity*, and denote it \mathcal{M}_g .

3.2 The radiation model

Some authors have criticized the gravity model, since it is unable to predict different fluxes between locations with similar densities and at the same distance, but have different population densities between them. They therefore proposed the radiation model, for which the expected number of transactions between two points depends on the number and size of other points around them. Formally, the number of expected transactions between v_i and v_j using the radiation model [18] is given by:

$$r_0(e_{ij}) = \frac{T_i}{1 - \frac{n_i}{N}} \times \frac{n_i \times n_j}{(n_i + s_{ij}) \times (n_i + n_j + s_{ij})}$$

with:

- s_{ij} the population in a circle whose center is v_i and radius d_{ij} minus n_i and n_j .
- $T_i = \sum_j W_*(e_{ij})$.
- N the total population.

While the original radiation model is parameter free, and thus do not have a deterrence function, a recent improvement [17] has been proposed. Using the same mechanism as the one used in gravity model to learn a deterrence function from data, it is defined as:

$$r(e_{ij}) = r_0(e_{ij}) \times f_2(d_{ij})$$

where $f_2(d_{ij})$ is the deterrence function defined as:

$$f_2(d) = \frac{\sum_{x,y|d_{xy}=d} W_{\star}(e_{xy})}{\sum_{x,y|d_{xy}=d} r(e_{xy})} \quad (3)$$

We will refer to this model as *radiation*, and denote it \mathcal{M}_r .

4 Experiments

4.1 Synthetic experiments

To test the effectiveness of the mobility models, we performed synthetic experiments. We adapted the benchmark used in [17] to our problem of pattern mining. The main idea is based on generating different synthetic datasets in which several patterns (exceptional subgraphs) are hidden, and evaluating the ability of the different models to identify these patterns. We picked up uniformly at random N vertices in the geographic space located in the bounding box between (45.8, 4.7) and (46.05, 4.95) (corresponding to the location of Lyon), and assigned to each vertex v_i a population n_i using one of following options: (1) *uniform*: we assign the same population $n_i = 100$ for all the vertices, (2) *random*: we assign a population uniformly at random from $[1, \dots, 100]$.

We injected 10 different patterns that constitute the set $P = \{p_1, \dots, p_{10}\}$. Each pattern p_i is made of 5 to 10 edges, and each edge has between 10 to 20 transactions. The first 5 patterns are associated to specific contexts that are different from \star ($C = (C_1, \dots, C_p)$ with $C_i \neq \star_i$ for nominal attributes, and $C_i \in [\min \mathbf{dom}(R_i), \max \mathbf{dom}(R_i)]$ for numerical attributes), and the 5 last ones are associated to $\star = (C_i, \dots, C_p)$ with $C_i = \star_i$ or $C_i = \mathbf{dom}(R_i)$. Then, noise is added to the transactions, governed by *noise_rate*. We generate a set of other transactions called *typical* transactions, that are distributed w.r.t a particular probability distribution $\mathcal{P}(t_{ij})$. These transactions represent the simple trajectories that do not depict an exceptional behavior. They represent 50% of the overall dataset and are generated using either (1) *distance*: $\mathcal{P}(t_{ij}) = \frac{n_i \times n_j}{d_{ij} \times Z_1}$. Where Z_1 is a normalization factor $Z_1 = \sum_{i,j} \frac{n_i \times n_j}{d_{ij}}$, or (2) *flux*: $\mathcal{P}(t_{ij}) = \frac{n_i \times n_j}{(n_i + s_{ij}) \times (n_i + n_j + s_{ij}) \times Z_2}$. Where Z_2 is a normalization factor $Z_2 = \sum_{i,j} \frac{n_i \times n_j}{(n_i + s_{ij}) \times (n_i + n_j + s_{ij})}$. Thus, we have 4 different configurations: (1) *distance uniform*, (2) *distance random*, (3) *flux uniform*, and (4) *flux random*.

For each generated dataset, we ran COSMIC with the three different models: \mathcal{M}_0 , \mathcal{M}_g and \mathcal{M}_r . We evaluated the result of each of them with the *Precision*, *Recall*, and *fScore* measures adapted to our case study. Let $P = \{p_1, \dots, p_n\}$ be the set of generated patterns, and $R = \{r_1, \dots, r_m\}$ the set of returned patterns by Cosmic. The precision of a pattern $r_i \in R$ w.r.t P is computed by: $Precision(r_i, P) = \max_{p_j \in P} \frac{|r_i \cap p_j|}{|r_i|}$.

The recall of a synthetic pattern $p_i \in P$ w.r.t R is: $Recall(p_i, R) = \max_{r_j \in R} \frac{|p_i \cap r_j|}{|p_i|}$.

Finally, the global precision, recall, and F-score of the two sets of patterns (P, R) is:

$$Precision(R, P) = \frac{\sum_{r_i \in R} Precision(r_i, P)}{|R|}, Recall(R, P) = \frac{\sum_{p_i \in P} Recall(p_i, R)}{|P|}, \text{ and } Fscore(R, P) = 2 \times \frac{Precision(R, P) \times Recall(R, P)}{Precision(R, P) + Recall(R, P)}.$$

Each configuration was executed 10 times.




Fig. 1 F-score, precision, and recall of each model in the different configurations.

Fig. 1 presents F-score, precision and recall of each model in each configuration, w.r.t different values of *noise_rate*. F-score and recall of mobility models are always

better than those of \mathcal{M}_0 whereas the precision remains similar. \mathcal{M}_r model is slightly better than \mathcal{M}_g when using the *flux* distribution probability. In fact, \mathcal{M}_r calculation is based on the *flux* hypothesis.

4.2 Experiments on real data

In this section, we compare the models \mathcal{M}_0 , \mathcal{M}_g and \mathcal{M}_r using a real world dataset. To this end, we first present the data and we show some statistics about them. Second, we study the distribution of transactions predicted by each model. Finally, we compare the result patterns detected using these different models.

We use the VÉLO'v dataset. VÉLO'v is the bike-sharing system run by the city of Lyon (France) and the company JCDecaux². There are a total of 348 VÉLO'v stations across the city of Lyon. Our experiments are performed on trips collected on October 2011. The overall number of trips is 565,065 transactions. Each trip includes the bicycle station and the time stamp for both departure and arrival, as well as some basic demographics about the users (gender, age, zip code, country of residence, type of pass). Hence, the VÉLO'v stations are the graph vertices ($|V| = 348$), and directed edges correspond to the fact that a VÉLO'v user checks out a bicycle at a station and returns it at another.

Fig. 2 (left) reports the distribution of populations of areas containing stations. This figure shows that the population is not uniform. Thus, it will inevitably influence the results of gravity and radiation models. Fig. 2 (right) shows the distribution of distance of transactions. There is only few transactions with distances less than 200 meters. In fact, people find it useless to take bikes for very short distances. The number of transactions increases when the distance increases until it reaches its maximum values for distances around 1km and 2km. After that, the number of transactions decreases.




Fig. 2 Statistics about the VÉLO'v dataset, (left) distribution populations of areas containing stations (right) distribution of distance of transactions.

² <http://www.velov.grandlyon.com/>

Fig. 3 and 4 compare expected edges starting from specific stations located respectively in Part Dieu and Cordeliers. They correspond to edges that contain at least 10 expected transactions. Recall that in \mathcal{M}_0 model, the expected number of transactions is simply the observed number of transactions. For example, if we take the station located in Part-Dieu, in \mathcal{M}_0 model some edges are connected to areas located in the north west even if they are far away from Part Dieu. This exceptional-ity will be captured by mobility models since these edges are not expected by them. Meanwhile, it is clear that the distribution of expected transactions strongly depend on the distance for spatial models.



Fig. 3 Comparison of expected edges starting from a station located in Part Dieu.



Fig. 4 Comparison of expected edges starting from a station located in Cordeliers.




Fig. 5 Distribution of expected transactions with respect to n_i (source population), n_j (destination population), d_{ij} (distance between source and destination), and WR_{acc} values, for each model.

Fig. 5 shows the distribution of expected transactions with respect to the following parameters: d_{ij} (distance between source and destination), n_i (source population), and n_j (destination population). Distributions of transactions by distance are similar for \mathcal{M}_0 and \mathcal{M}_g , while \mathcal{M}_r expects much more for distances less than 1km. For the two other distributions, \mathcal{M}_r expects more than the other models in most of the cases. Also, the number of expected transactions by \mathcal{M}_0 is higher than \mathcal{M}_g when n_i and n_j are less than 3000, but it is lower when n_i and n_j exceed 3000. In fact, \mathcal{M}_g expects more transactions when n_i and n_j are greater.

We computed the result patterns and we compared their ranking with respect to the three different models. To this end, we ranked the patterns based on their *WRAcc* scores using each model, and we computed the kendall tau between these rankings. Results are depicted in Fig. 6. As expected, the greatest kendall tau is between \mathcal{M}_g and \mathcal{M}_r , which is logical since they are spacial models based on similar information. The kendall tau of \mathcal{M}_0 with the spacial models is lower, especially with \mathcal{M}_r .

	\mathcal{M}_0	\mathcal{M}_g	\mathcal{M}_r
\mathcal{M}_0	100 %	69 %	63 %
\mathcal{M}_g	69 %	100 %	79 %
\mathcal{M}_r	63 %	79 %	100 %

Fig. 6 Kendall tau of ranking of results based on the three different models \mathcal{M}_0 , \mathcal{M}_g and \mathcal{M}_r .




Fig. 7 Boxplots of $avg(\frac{n_i \times n_j}{d_{ij}})$ of exceptional transactions of each model.

Fig. 8 reports the top 5 patterns for each model. Gravity and radiation models have found exactly the same 5 best patterns, whereas \mathcal{M}_0 model differs in some of them. In order to verify whether the spacial models have a real impact in the top patterns, we have compared the average value of $\frac{n_i \times n_j}{d_{ij}}$ of the transactions that appear in the top 10 patterns in each model. Fig. 7 presents the results. This value is significantly lower in the spacial models comparing with the \mathcal{M}_0 one. It means that the spacial models overweight the *WRAcc* quality of transactions with low values of $\frac{n_i \times n_j}{d_{ij}}$.

5 Conclusion

In this paper, we extended the problem of discovering exceptional contextual sub-graphs by considering two well-known models from statistical physics: The gravity and the radiation models. Experiments on both synthetic and real-world datasets

Fig. 8 Top 5 patterns found by each model \mathcal{M}_0 , \mathcal{M}_g and \mathcal{M}_r .

demonstrate that considering demographic and spatial information makes it possible to capture more complex phenomena.

Acknowledgement. This work has been partially supported by the projects GRAISearch (FP7-PEOPLE-2013-IAPP), VEL'INNOV (ANR INOV 2012) and the Group Image Mining (GIM) which joins researchers of THALES Group and LIRIS Lab. We thank especially Jérôme Kodjabachian and Bertrand Duqueroie of AS&BSIM Lab.

References

1. Ashbrook, D., Starner, T.: Using gps to learn significant locations and predict movement across multiple users. *Pers. Ub. Comput.* **7**(5), 275–286 (2003)
2. Atzmueller, M.: Detecting community patterns capturing exceptional link trails. In: 2016 IEEE/ACM ASONAM, pp. 757–764 (2016)
3. Bayir, M.A., Demirbas, M., Cosar, A.: Track me! a web based location tracking and analysis system. In: ISICIS, pp. 117–122. IEEE (2009)
4. Cazabet, R., Borgnat, P., Jensen, P.: Enhancing space-aware community detection using degree constrained spatial null model. In: Workshop CompleNet, pp. 47–55. Springer (2017)
5. Demirbas, M., Rudra, C., Rudra, A., Bayir, M.A.: imap: Indirect measurement of air pollution with cellphones. In: PerCom Workshops, pp. 1–6 (2009)
6. Expert, P., Evans, T.S., Blondel, V.D., Lambiotte, R.: Uncovering space-independent communities in spatial networks. *PNAS* **108**(19), 7663–7668 (2011)
7. Giannotti, F., Pedreschi, D.: Mobility, data mining and privacy. Springer Science & Business Media (2008)
8. Harrington, A., Cahill, V.: Route profiling: putting context to work. In: SAC, pp. 1567–1573 (2004)
9. Kaytoue, M., Plantevit, M., Zimmermann, A., Bendimerad, A.A., Robardet, C.: Exceptional contextual subgraph mining. *Machine Learning* **106**(8), 1171–1211 (2017)
10. Lavrac, N., Flach, P., Zupan, B.: Rule evaluation measures: A unifying view. In: ILP-99, pp. 174–185 (1999)
11. Leman, D., Feelders, A., Knobbe, A.J.: Exceptional model mining. In: ECML/PKDD, pp. 1–16 (2008)
12. Lemmerich, F., Becker, M., Singer, P., Helic, D., Hotho, A., Strohmaier, M.: Mining subgroups with exceptional transition behavior. In: ACM SIGKDD, pp. 965–974 (2016)
13. Li, Z., Ji, M., Lee, J.G., Tang, L.A., Yu, Y., Han, J., Kays, R.: Movemine: Mining moving object databases. In: SIGMOD, pp. 1203–1206. ACM (2010)
14. Luo, W., Tan, H., Chen, L., Ni, L.M.: Finding time period-based most frequent path in big trajectory data. In: SIGMOD, pp. 713–724. ACM (2013)
15. Masucci, A.P., Serras, J., Johansson, A., Batty, M.: Gravity versus radiation models. *Physical Review E* **88**(2), 022,812 (2013)
16. Monreale, A., Pinelli, F., Trasarti, R., Giannotti, F.: Wherenext: A location predictor on trajectory pattern mining. In: KDD, pp. 637–646. ACM (2009)
17. Sarzynska, M., Leicht, E., Chowell, G., Porter, M.: Null models for community detection in spatially embedded, temporal networks. *J. Complex Networks* **4**(3), 363–406 (2016)
18. Simini, F., González, M.C., Maritan, A., Barabási, A.L.: A universal model for mobility and migration patterns. *arXiv preprint arXiv:1111.0586* (2011)
19. Wang, D., Pedreschi, D., Song, C., Giannotti, F., Barabasi, A.L.: Human mobility, social ties, and link prediction. In: KDD, pp. 1100–1108. ACM (2011)
20. Zheng, Y., Zhang, L., Xie, X., Ma, W.Y.: Mining interesting locations and travel sequences from gps trajectories. In: WWW, pp. 791–800. ACM (2009)
21. Zipf, G.K.: The p 1 p 2/d hypothesis: The case of railway express. *The Journal of Psychology* **22**(1), 3–8 (1946)