Mining Exceptional Social Behaviour

Carolina Centeio Jorge[1,2], Martin Atzmueller[3], Behzad M. Heravi[4],
Jenny L. Gibson[5], Cláudio Rebelo de Sá[2], and Rosaldo J. F. Rossetti[1]

[1] University of Porto
[2] University of Twente
[3] University of Tilburg
[4] University College London
[5] University of Cambridge

Abstract. Essentially, our lives are made of social interactions. These can be recorded through personal gadgets as well as sensors adequately attached to people for research purposes. In particular, such sensors may record real time location of people. This location data can then be used to infer interactions, which may be translated into behavioural patterns. In this paper, we focus on the automatic discovery of exceptional social behaviour from spatio-temporal data. For that, we propose a method for Exceptional Behaviour Discovery (EBD). The proposed method combines Subgroup Discovery and Network Science techniques for finding social behaviour that deviates from the norm. In particular, it transforms movement and demographic data into attributed social interaction networks, and returns descriptive subgroups. We applied the proposed method on two real datasets containing location data from children playing in the school playground. Our results indicate that this is a valid approach which is able to obtain meaningful knowledge from the data.

Keywords: Subgroup Discovery · Network Science · Social Interactions.

1 Introduction

People interact everyday through verbal and non-verbal communication. These interactions allow us to study human beings as social entities [10]. From that, some phenomenons may emerge [6], such as homophily [25], the tendency of people for interacting more with those who are more similar to them, and so forth. This suggests that socio-demographic characteristics, as well as behavioural patterns, tend to be localized [25]. Thus, the automatic detection of patterns and unusual behaviour can be valuable to the understanding and discovery of the interactions in marketing [11], education [1], security [22], and health [28].

Furthermore, as people make more and more use of new technology, a great amount of data from their behaviour is being collected [33]. In addition, deliberately gathering data from social and ubiquitous environments through sensors (of proximity or geo-localization) is also being used to study the behaviour of people without interfering with their actions [16], e. g., *movement data* [23], on a more naturalistic basis. More specifically, it was applied to the domain of social

interactions [26], where the authors analyze the properties (mainly focusing on gender phenomena) of five-year-old kids' social interactions.

Interactions may follow patterns, sequences of behaviours, or be expressed with verbal and non-verbal gestures which we do not even notice [14]. In particular, there may be some patterns which do not follow the norm, making them unusual, leading to *Exceptional Behaviour Discovery* (EBD). In that context, the automatic extraction of descriptive knowledge from the data, such as subgroups, can help and support the analysis and decisions of social sciences experts.

We propose a method for *Exceptional Behaviour Discovery* which is a combination of Data Mining approaches. In particular, we propose the combination of Subgroup Discovery and Network Science methods for the automatic detection of the characteristics that can better describe unusual social interactions. The main goal of the proposed method is to find subgroups from movement data, making use of a graph structure. Thus, we focus on descriptive subgroup discovery on such (extended) graph data structures. On the one hand, Subgroup Discovery [20] is a descriptive data mining technique that provides easy-to-understand results to the expert. It finds subgroups of objects in the data that share the same characteristics with respect to a property of interest (target) [17]. On the other hand, interactions can be represented as a set of complex networks, namely *social interaction networks* which capture interactions between the subjects involved in the study [4]. In this case, people are represented by nodes and the interactions are represented by edges.

When considering social interaction data, however, people interact on the move and over time; therefore, the sequence of locations of a person can be related to the demographic properties of the subjects involved, and modeled as attributed social interaction networks. Eventual behavioural patterns thus might not be captured in snapshots of the network, but rather in the evolution of it. Building upon compositional subgroup discovery on such attributed social interaction networks [6], we created different Subgroup Discovery quality measures to find subgroups of people whose interactions deviate from the norm.

We tested the proposed approaches on two sets of data with locations and personal attributes of children in the playground of the school. The data was collected with the use of location sensors during the school breaks. One dataset, *playgroundA* [16], has the geographic position of 18 children over time, in 10 different sessions and personal attributes (gender, age, emotional stability etc). The other dataset, *playgroundB* [26] has the position of 16 children and sociodemographic attributes, such as gender and age. The results were mostly expected, when analyzed by experts [26] in the domain, and similar between the two datasets. This shows it is a valid approach. Also, the results added meaningful information to the expected scenario.

The remainder of this paper is structured as follows: in Section 2 we present the background, in which we explain the underlying concepts and literature review; in Section 3 we present our contributions for the state of the art, testing it with a case study presented in Section 4; we finally conclude in Section 5.

2   Preliminaries

Many domains in which we can potentially use data mining techniques are placed in a temporal or spatial scenario. Therefore, to learn from the data, it is important to take into account its temporal and spatial properties [29]. With spatial properties providing information about objects' location, also known as *movement data* [23].

2.1   Subgroup Discovery

Subgroup Discovery (SD) is a descriptive and exploratory data mining technique to identify interesting patterns, the so-called subgroups, that deviate from the norm [20]. These patterns show an unusual distribution when compared to the overall population [3]. This interesting behaviour is typically based on some criteria which balances their relevance between their size and unusualness. We can find SD applications on medical [13], marketing [8], education [30], socio-demographic [21] and social domains [6].

As in [12], we define a dataset as a bag of n records given in the form of $x = (a_1, \ldots, a_m, t_1, \ldots, t_l)$, where a_i is a descriptor and t_i is a target. Subgroups are usually described with a description language, $\mathcal{D}$, and are induced by a *pattern*. A *pattern*, p, is a function $p : \mathcal{A} \rightarrow \{0,1\}$ and *covers* a record x iff $p(a_1, \ldots, a_m) = 1$. A *subgroup* corresponding to a pattern p is the bag of records, S_p, that p covers: $S_p = \{x \in D \mid p(a_1, \ldots, a_m) = 1\}$. $\mathcal{D}$ is typically a conjunction of conditions on attributes, such as: Gender $=$ F $\wedge$ Age ≤ 22.

The interestingness of subgroups is measured by *quality measures* according to the different types of targets. Given a subgroup discovery algorithm, a set of subgroups is identified and scored by the quality function [24]: $\varphi : \mathcal{D} \rightarrow \mathbb{R}$. Quality measures are a key factor for the extraction of knowledge because the interest obtained depends directly on them [17]. Many have been proposed for identifying different deviations in different targets. Targets can be binary [35], nominal [8], numeric [15], ranked [31] or as a distribution [18].

2.2   Network Science

Network Science combines ideas from several domains of knowledge so as to address questions about networks [27]. A network is a collection of *nodes* connected with *edges*. This simple representation allows one to translate many events into the form of networks, which can often lead to new and useful insights [27].

The key concepts of Network Science are centrality measures, which measure the nodes that are the most important or central in a network. Centrality gathers a wide range of metrics and measures that can allow us to better understand the data. For example, *degree* of centrality (based on the number of links of a node), *closeness* (based on the average length of the shortest path between the node and all other nodes in the graph), *betweeness* (based on how many shortest paths of the graph go through a node) and *pagerank* (measured by the links to a node). More recently proposed metrics are *hubs* and *authorities* [19]. A hub is

a node with many outgoing links to authorities, whereas an authority is a node with many links from hubs. Another network concept of practical importance is provided by communities [27] in networks. Communities are tightly knit groups within a larger, looser network.

A particular case of networks are social interaction networks [34] which focus on interactions between people as the corresponding actors. In this case, the nodes represent the actors and the edges, the links between actors, model a interaction or event. These edges may have properties, such as frequency of occurrence or duration. Furthermore, edges and nodes may have other labels, leading to attributed networks. From these attributed networks, we can extract and characterize subgroups [6].

A complex network can be represented by a graph [9]. A graph G is an ordered triple $(V(G), E(G), \psi_G)$, where $V(G)$ represents the set of vertices, $E(G)$, the edges and ψ_G is the function that associates to each edge of G a pair of vertices of $V(G)$. For example: $V(G) = \{v_1, v_2, \ldots, v_n\}$, $E(G) = \{e_1, e_2, \ldots, e_n\}$ and $\psi_G(e_1) = (v_1, v_2)$. A graph can be *directed* or *undirected*. In the case of G being directed, the output of the function $\psi_G(e_i)$, (v_j, v_k) is ordered and it is known as a *digraph* [27]. Moreover, the graph can have multiple edges, in the same direction, if directed, between two nodes. In this case, the graph is referred to as *multigraph*. The function ψ_{MG} of a multigraph returns the same pair of vertices for more than one edge.

Some approaches combine Subgroup Discovery and Network Science. In 2013, Atzmueller [2] gave an overview of data mining in social interaction networks, specifically human behavioural (offline) networks. Methods and approaches for describing and characterizing networks and their properties were proposed. In terms of community detection, Skrlj et al. [32] introduced the Community-Based Semantic Subgroup Discovery (CBSSD), an algorithm that identifies classes of instances based on structural properties of complex networks. Atzmueller [5] also presented an overview of research in subgroup discovery and community detection on attributed graphs. In 2018, Atzmueller [6] also proposed quality measures and targets based on interaction network properties for subgroup discovery in attributed social networks, as compositional subgroup discovery.

3   Exceptional Behaviour Discovery

In this paper we propose *Exceptional behaviour discovery (EBD)*. The aim of EBD is to look for social behaviour which deviates from the norm. In order to recognize unusual social behaviour among individuals (in social interactions), we adapted an existing subgroup discovery technique to deal with spatio-temporal data. We focus on the study of Subgroup Discovery methods and metrics of social networks analysis. This work extends the work proposed in [6] which combined Subgroup Discovery with social interaction networks and is referred as *Compositional Subgroup Discovery*.

3.1 Compositional Subgroup Discovery

Compositional Subgroup Discovery can be divided into two steps. First, the network is represented by means of a graph, where each subject is represented by a node and each interaction is represented by an edge between two nodes. In this graph representation, both nodes and edges can be characterized by attributes. Finally, these can be used to find subgroups and to explain some observed behavioural patterns.

Quality measures To measure the interestingness, the duration of the interactions and frequency are considered. The target, t_p, is numeric and corresponds to the observed number of edges normalized by the expectation.

Two different quality measures were proposed. The first measure uses simple attributed graphs where the duration of the interactions is used to weight the edges. The second one also includes the interaction frequency information in an attributed multigraph representation of the network (each edge represents an interaction).

In the first approach, the simple attributed graph, the weights of all the edges, E_p, covered by a *pattern*, are summed, normalized by the number of possible edges, n_E, among the nodes covered by the *pattern*, n_{Ep}. Then, r samples of n_E edges, where $n_E = \frac{n_{Ep}(n_{Ep}-1)}{2}$, are considered as well as are their normalized sum of weights. Finally, a Z-score is calculated estimating the significance of the obtained value (t_P) among the samples. For a *pattern p*, the quality function, q_S, is:

$$q_S(P) = Z(\frac{1}{n_E} \cdot \sum_{e \in E_P} w(e)) \tag{1}$$

In the multigraph version, the frequency (apart from the duration) of interaction is also taken into account. Thus, for normalizing the sum of weights of a pattern p, we have to consider the multiple edges that exist between two nodes. In this case, instead of dividing by n_E, the author divides by $n_e + m_E$, where $m_E = \sum_{i=1}^{n_E}(m_i - 1)$ and m_i is the observed multiplicity of an edge. For the Z-score, all the edges are considered.

3.2 Proposed Extension

In this work, we propose to use *digraphs* to represent the interactions of the subjects. For that, we need to define proximity and when a subject is approaching another subject. If a subject approaches another within a certain proximity, a directed edge is created from the node of the subject approaching to the node of the subject approached.

This approach combines *movement data* and *social data* of the subjects and returns subgroups, according to the desired quality function. The *movement data* consists of a timestamp of the event, the *id* and position (x and y) for each subject. From that, there is a function that computes the speed, *velX* and

velY, relatively to x and y, respectively. The *social data*, has the ids and socio-demographic data corresponding to the subjects in the *movement data*. Any numeric attributes are discretized in equal frequency bins.

Generating the interaction digraph To create the interaction digraph, or multidigraph, we first need to define interactions. We consider an interaction between two subjects when their relative distance is within a certain proximity and one of the subjects approaches the other. Therefore, given a maximum distance threshold between subjects, *maxdist*, we start with an empty digraph G.

At each time step t, a matrix of distances, D, between every two subjects is calculated. Then, for each distance $d_{i,j} \in D : d_{i,j} \leq maxdist$ we compute a vector from i to j as $\boldsymbol{r}_{i,j} = (x_j - x_i, y_j - y_i)$. We then verify the speed vector of i, $\boldsymbol{vel}_i = (velX_i, velY_i)$, and calculate the cosine between the vectors $\boldsymbol{r}_{i,j}$ and $\boldsymbol{vel}_i$. If the cosine be positive, we consider that the subject i *approached* (or *reached*) subject j.

In the simple digraph, a directed edge, from node i to node j, is added to G and $w_{i,j} \in W$ is incremented one unit of time, where W is the matrix of weights and $w_{i,j}$ is the number of times that the subject i approaches the subject j. In the multidigraph version, a directed edge is added to G at moment t if subject i approached subject j, given that it was not interacting in $t - 1$, with $w_{i,j} \in W$, where $w_{i,j}$ is the total time that the subject i approaches the subject j without interruption.

Quality measures We propose two quality measures with two variations.

Simple attributed digraph. This quality measure takes into account the duration of the interactions between two subjects. A new directed edge (or arrow) is considered every time an interaction is observed and not clear.

For the quality function, we use the same measure as q_S (Equation 1). However, since we have the double of the edges (because this is a directed version), we use $n_E = n_{Ep}(n_{Ep} - 1)$.

Directed attributed multidigraph. This quality measure considers both the duration and frequency of the interactions between two subjects. In this case, one directed edge is created every time an interaction is observed.

To-node and From-node *variants.* These two variants, *To-node* and *From-node*, extend the quality measures mentioned above. In the *To-node* and *From-node*, the attributes of the edges are only based on the attributes of the *head* node or the *tail* node, respectively. With these variants we hope to find valuable information about the attributes of the subjects that look for interactions (*From-node*) and the subjects that are reached the most (*To-node*).

3.3   Subgroup Discovery

The edges of the graph G are associated with features, which are based on the attributes of the nodes of that edge. Numeric attributes in the nodes are represented as *equal* (or same), *greater* or *lower* in the features of the edges in the comparison versions (Simple attributed digraph and Directed attributed multidigraph). In the *To-node* and *From-node* variants, the numeric attributes of the nodes are represented as *equal* (or same), *high* or *low* as the features of the edges. After assigning attributes to the edges, there is a function to run the SD-Map algorithm [7] on the graph. The output is a list of subgroups and their characteristics, namely pattern description, the number of edges and nodes covered by the pattern, the mean weight of those edges and the score (quality function result).

We also propose to add automatically generated features to the nodes' attributes and, consequently, to the edges' with the use of complex networks' metrics. For example, degree (also in-degree and out-degree), centrality measures (eigen, closeness, betweenness), authority and hub values.

4   A Case Study in Playground Social Interactions

Analyzing social interactions in the playground can be of utmost importance. Social group structure and dynamics are believed to be strongly related to the child well-being and yet has been poorly understood and studied [16].

4.1   Data

To test the approach explained in Section 3.2, we used two datasets with locations of children in the school playground. The data was collected with the use of location sensors during the school breaks.

One dataset, *playgroundA* [16], has the geographic position of 18 7-8 year-old children (9 girls) over time, during approximately 45min. It also includes the personal attributes (gender, age, emotional stability etc). The kids were playing outdoors, without toys, during a normal day of Primary School. They had a head-mounted sensor with IMU and GNSS for precise positioning a shoe-mounted IMU sensor for activity monitoring. The following social and psychological measures were collected from a teacher:

ProSoc the higher the score the highest the social skills are
Conduct a high value represents behaviour problems
Emotion the higher the score the more emotional difficulties
Peer high scores indicate that the child has issues making friends
Hyper the higher the more hyperactive the child is

The other dataset, *playgroundB* [26] has the position of 14 children (8 girls) around 5 years-old and socio-demographic attributes, such as gender and age. The data was collected through a real time location system that used UWB sensors. The data was collected during 1h.

4.2   Assessment Approach

The first dataset already had ids, positions, and speeds in two dimensions. For the *social data* of this dataset, we transformed the numeric values in 3 bins. The second dataset did not include speed values, so we created them. Then, we created the graphs and experimented the 6 approaches for each: comparison, to-node and from-node attributed edges for both simple and multidigraph versions. Furthermore, we also looked for subgroups based on Network Science metrics in the *playgroundB* dataset.

4.3   Results and Discussion

In this section, we analyze some of the results of our approach with both *playgroundA* and *playgroundB* datasets. An adapted version of SD-Map [7] is used for subgroup discovery.

PlaygroundA Table 1 shows the ranked subgroups found with the dataset *playgroundA*. We present three versions: a comparison version with simple attributed digraphs (comp in the column V, in the Table 1), and a to-node and from-node version with attributed multidigraph version (to and from in the column V, in the Table 1, respectively). For each subgroup, we show its pattern, number of nodes (children) belonging to the subgroup, N, number of edges (interactions), E, the mean time of an interactions between children in the subgroup, $|C|$, and the Z-score based on the comparison between the total duration of the interactions in the subgroup and the null model, Z.

The top-2 ranked subgroups (Table 1) obtained with the comparison version are $Gender = M \rightarrow Gender = M$ and $Gender = F \rightarrow Gender = F$. This means that boys follow other boys and girls follow other girls, much more than what would be expected. We note that these subgroups have much higher score than the others, which goes in line with the observations already discussed in [26]. This seems to confirm the homophily regarding the gender, meaning that children interact more with children of the same gender.

The subgroup $Emotion = higher \rightarrow Emotion = lower$ seems to indicate that children with more emotional problems tend to look for interactions with children with less emotional issues. The 4th best subgroup $Hyper = same \rightarrow Hyper = same$ suggests that children prefer to look for other children with similar levels of peer relations. This is also indicated by the 8th subgroup, since it connotes that similar social skills can also motivate more interactions.

Some subgroups exhibit opposite behaviour. For example, the subgroups 3 and 7 from the comparison in Table 1. However, when we visualize the graphs associated with the subgroups in Fig. 2 we can conclude that they are representing distinct behaviours. The shades in blue represent the weight of the edge: the lighter the blue the lower the weight. It is possible to observe that the subgroup ranked higher, in position 3, "Emotion=higher → Emotion=lower", presents edges with bigger weights. This makes sense since the quality function is based

on the sum of the weights of the edges. In other words, a subgroup is more interesting if the sum of the weights of its graph are bigger than the graph of all the interactions.

Results of the to-node and from-node versions can add valuable information to the results found in the comparison version. For these results, we conclude that the multidigraph version presents better results than the simple version. The results show that there is a tendency for older children to go after interactions, whereas average aged children are the ones reached the most. Children with a low 'Peer' score, meaning they present a better quality in peer relation, as well as low 'Hyper' and 'Emotion' scores, which means they do not present hyperactivity or emotional issues, both look for interactions and are reached by other peers. Furthermore, children with low social skills ("ProSoc") tend to reach for interactions whereas children with average social skills are more reached.

Table 1: Ranking of subgroups (comp, to-node and from-node attributed multidigraph versions) according to the total duration of interactions between every two children in the dataset *playgroundA*.

Rank	V	Pattern	N	E	$|C|$	Z
1	comp	Gender=M → Gender=M	9	51	21.1	28.6
2	comp	Gender=F → Gender=F	9	50	15.4	19.5
3	comp	Emotion=higher → Emotion=lower	18	73	7.4	3.9
4	comp	Hyper=same → Hyper=same	18	72	7.3	3.7
5	comp	Conduct=lower → Conduct=higher	18	74	7.2	3.4
6	comp	Age=higher → Age=lower	18	85	7.9	3.2
7	comp	Emotion=lower → Emotion=higher	18	74	7.0	3.1
8	comp	ProSoc=same → ProSoc=same	18	69	6.5	2.8
9	comp	Conduct=higher → Conduct=lower	18	75	6.8	2.8
10	comp	Peer=same → Peer=same	18	105	9.0	2.8
1	to	Conduct=low ∧ Peer=low ∧ Hyper=low	18	174	1.5	2.0
2	to	Age=Medium ∧ ProSoc=Medium ∧ Emotion=low	17	184	1.6	2.0
3	to	Age=Medium ∧ Emotion=low	17	184	1.6	1.9
4	to	Age=Medium ∧ Conduct=low	16	157	1.4	1.7
1	from	Peer=low ∧ Age=high ∧ Hyper=low	18	135	1.3	2.8
2	from	Peer=low ∧ Emotion=low ∧ ProSoc=low	18	158	1.4	2.4
3	from	Age=high ∧ Hyper=low	18	135	1.3	2.2
4	from	Gender = M ∧ Emotion=low ∧ Hyper=low	18	147	1.3	2.2

PlaygroundB For the dataset *playgroundB*, we analyze the attributed multidigraph approach. When analyzing the results we can also conclude that children in this dataset interact more with peers of the same gender. Moreover, we can see that boys tend to look for interactions with older boys, whereas girls show more interactions with girls with the same age.

If we focus on the to-node version of *playgroundB*, we can see that boys (Gender=M) are the most reached, regardless of their age. Nevertheless, the pattern with the highest score is "Gender=M ∧ Age=low". The oldest children, however, are the ones looking for more interactions according to the from-node

Emotion=higher → Emotion=lower   Emotion=lower → Emotion=higher

Fig. 2: Plots of graph representation of the subgroups 3 and 7, comparison version, in Table 1.

multidigraph version. In this version, all top-3 patterns include "Age=high", despite the gender, with small differences in the scores (11.9, 10.0 and 9.7).

Since we only have two attributes in this dataset (gender and age) we generated extra features based on the networks' metrics (Section 3.3). The results of the comparison version of simple attributed digraph are presented in Table 2. We can observe that boys tend to look for boys with a similar hub score and that girls look for girls with similar closeness. We can associate the hub score to interactions with popular kids and conclude that boys prefer to interact with boys with a similar level of interactions with popular peers. Closeness, on the other hand, may imply many interactions in general, which suggests that girls prefer to interact with girls that interact with a similar amount of peers. In general, we observed that children reach peers with similar centrality measures.

Table 2: Top-4 subgroups (comparison version of simple attributed digraph) according to the total duration of interactions between every two children, considering Network Science metrics in the dataset *playgroundB*.

Rank	Pattern	N	E	$\|C\|$	Z
1	Gender=M → Gender=M	6	25	0.2	21.7
2	Gender=M ∧ hubs=same → Gender=M ∧ hubs=same	4	12	0.2	11.7
3	Gender=F ∧ closeness=same → Gender=F ∧ closeness=same	6	12	0.1	7.7
4	Gender=F → Gender=F	8	22	0.1	7.0

5   Conclusions

In this paper we proposed an approach to extract descriptive knowledge about exceptional behaviour from demographic and movement data of social interactions. We extended an existing approach which combines Subgroup Discovery and Network Science techniques to find subgroups in attributed digraphs. Our main contributions include adapting this approach to movement data (data that represents location over time) and, as such, to directed digraphs, as well as adding Network Science metrics to the attributes of the graph. Accordingly, we

developed a pipeline that receives spatio-temporal data of tracked objects (people) along with some personal and social characteristics of the individuals. Then it transforms the data into attributed directed digraphs (simple and/or multidigraphs) and returns subgroups. To test our approach we used two datasets of children interacting in the school playground. The results were as expected by the experts in the domain and similar in both datasets. Nevertheless, they can add some valuable information for further social interaction analysis.

For future work, an interesting direction is given by further alternative quality measures that might be more refined to specific interaction contexts regarding the detection of subgroups of interactions. Furthermore, we also aim to compare the presented method with further approaches and to apply more datasets.

References

1. Altman, I.: The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding. Brooks/Cole Publishing Company (1975)
2. Atzmueller, M.: Data Mining on Social Interaction Networks. JDMDH **1** (2014)
3. Atzmueller, M.: Subgroup discovery. WIREs DMKD **5**(1), 35–49 (2015)
4. Atzmueller, M.: Local exceptionality detection on social interaction networks. In: ECML/PKDD (3). LNCS, vol. 9853, pp. 298–302. Springer (2016)
5. Atzmueller, M.: Descriptive community detection. In: Formal Concept Analysis of Social Networks, pp. 41–58. LNCS, Springer (2017)
6. Atzmueller, M.: Compositional subgroup discovery on attributed social interaction networks. In: DS. LNCS, vol. 11198, pp. 259–275. Springer (2018)
7. Atzmueller, M., Puppe, F.: SD-Map - A fast algorithm for exhaustive subgroup discovery. In: Proc. PKDD. LNCS, vol. 4213, pp. 6–17. Springer (2006)
8. Berlanga, F., del Jesus, M.J., González, P., Herrera, F., Mesonero, M.: Multiobjective evolutionary induction of subgroup discovery fuzzy rules: A case study in marketing. In: Perner, P. (ed.) Advances in Data Mining. Applications in Medicine, Web Mining, Marketing, Image and Signal Mining. pp. 337–349. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
9. Bondy, J.A., Murty, U.S.R.: Graph theory with applications, vol. 290. Elsevier Science Ltd/North-Holland (1976)
10. Cabrera-Quiros, L., Demetriou, A., Gedik, E., v. d. Meij, L., Hung, H.: The matchnmingle dataset: a novel multi-sensor resource for the analysis of social interactions and group dynamics in-the-wild during free-standing conversations and speed dates. IEEE Transactions on Affective Computing (2018)
11. Delener, N.: Religious contrasts in consumer decision behaviour patterns: Their dimensions and marketing implications. European Journal of Marketing **28**(5), 36–53 (1994)
12. Duivesteijn, W., Knobbe, A.J.: Exploiting false discoveries - statistical validation of patterns and quality measures in subgroup discovery. In: ICDM. pp. 151–160. IEEE Computer Society (2011)
13. Gamberger, D., Lavrac, N.: Expert-guided subgroup discovery: Methodology and application. J. Artif. Intell. Res. **17**, 501–527 (2002)
14. Goffman, E. (ed.): Interaction Ritual: Essays in Face to Face Behavior. Aldine-Transactions (1967)
15. Grosskreutz, H., Rüping, S.: On subgroup discovery in numerical domains. In: Proc. ECML/PKDD (1). LNCS, vol. 5781, p. 30. Springer (2009)

16. Heravi, B.M., Gibson, J.L., Hailes, S., Skuse, D.: Playground social interaction analysis using bespoke wearable sensors for tracking and motion capture. In: Proceedings of the 5th International Conference on Movement and Computing. pp. 21:1–21:8. MOCO '18, ACM (2018)
17. Herrera, F., Carmona, C.J., González, P., del Jesús, M.J.: An overview on subgroup discovery: foundations and applications. Knowl. Inf. Syst. **29**(3), 495–525 (2011)
18. Jorge, A.M., Pereira, F., Azevedo, P.J.: Visual interactive subgroup discovery with numerical properties of interest. In: Discovery Science. LNCS, vol. 4265, pp. 301–305. Springer (2006)
19. Kleinberg, J.M.: Hubs, authorities, and communities. ACM Comput. Surv. **31**(4es), 5 (1999)
20. Klösgen, W.: Handbook of Data Mining and Knowledge Discovery. Oxford University Press, Inc., New York, NY, USA (2002)
21. Klösgen, W., May, M.: Spatial subgroup mining integrated in an object-relational spatial database. In: Proc. PKDD. LNCS, vol. 2431, pp. 275–286. Springer (2002)
22. Ko, T.: A survey on behavior analysis in video surveillance for homeland security applications. In: 37th IEEE Applied Imagery Pattern Recognition Workshop, AIPR 2008, Washington, DC, USA, 15-17 October 2008, Proceedings. pp. 1–8 (2008)
23. Lauw, H.W., Lim, E., Pang, H., Tan, T.: Stevent: Spatio-temporal event model for social network discovery. ACM Trans. Inf. Syst. **28**(3), 15:1–15:32 (2010)
24. Leman, D., Feelders, A., Knobbe, A.J.: Exceptional model mining. In: ECML/PKDD (2). LNCS, vol. 5212, pp. 1–16. Springer (2008)
25. McPherson, M., Smith-Lovin, L., Cook, J.M.: Birds of a feather: Homophily in social networks. Annual Review of Sociology **27**(1), 415–444 (2001)
26. Messinger, D.S., Prince, E.B., Zheng, M., Martin, K., Mitsven, S.G., Huang, S., Stlzel, T., Johnson, N., Rudolph, U., Perry, L.K., Laursen, B., Song, C.: Continuous measurement of dynamic classroom social interactions. International Journal of Behavioral Development **43**(3), 263–270 (2019)
27. Newman, M.: Networks: An Introduction, chap. Introduction: A short introduction to networks and why we study them. Oxford University Press, Inc. (2010)
28. Owen, N., Healy, G.N., Matthews, C.E., Dunstan, D.W.: Too much sitting: the population health science of sedentary behavior. Exercise and sport sciences reviews **38(3)**, 182–196 (2010)
29. Roddick, J.F., Spiliopoulou, M.: A bibliography of temporal, spatial and spatio-temporal data mining research. SIGKDD Explorations **1**(1), 34–38 (1999)
30. Romero, C., González, P., Ventura, S., del Jesús, M.J., Herrera, F.: Evolutionary algorithms for subgroup discovery in e-learning: A practical application using moodle data. Expert Syst. Appl. **36**(2), 1632–1644 (2009)
31. de Sá, C.R., Duivesteijn, W., Soares, C., Knobbe, A.J.: Exceptional preferences mining. In: Proc. Discovery Science. pp. 3–18 (2016)
32. Skrlj, B., Kralj, J., Vavpetic, A., Lavrac, N.: Community-based semantic subgroup discovery. In: Proc. NFMCP@PKDD/ECML. LNCS, vol. 10785, pp. 182–196. Springer (2017)
33. Terry, M.A., Mynatt, E.D., Ryall, K., Leigh, D.: Social net: using patterns of physical proximity over time to infer shared interests. In: CHI Extended Abstracts. pp. 816–817. ACM (2002)
34. Wasserman, S., Faust, K.: Social Network Analysis: Methods and Applications. Structural Analysis in the Social Sciences, Cambridge University Press (1994)
35. Wrobel, S.: An algorithm for multi-relational discovery of subgroups. In: PKDD. LNCS, vol. 1263, pp. 78–87. Springer (1997)