

Descriptive Community Detection

Martin Atzmueller

1 Introduction

Subgroup discovery [6, 35, 67] aims at identifying interesting descriptive subgroups contained in a dataset—from a compositional network analysis view, aiming at a description given, e.g., by a set of attribute values. The subgroups are identified in such a way that they are interesting with respect to a certain target property. In the context of ubiquitous data and social media, interesting target concepts are given, e.g., by binary variables for obtaining characteristic descriptions of certain phenomena, densely connected graph structures (communities), or exceptional spatio-semantic distributions [4, 8]. This directly bridges the gap to community detection methods [26, 53, 69] that focus on structural aspects of a network/graph, for finding densely connected subgroups of nodes.

This paper, an extended and significantly revised version of [5] presents an organized picture of recent research in subgroup discovery and community detection specifically focusing on attributed graphs. We start with the introduction of necessary background concepts in Sect. 2. After that, we provide a compact overview on prominent methods for community detection, and discuss the exceptional model mining approach. Next, Sect. 3 describes recent work on mining attributed graphs for description-oriented approaches. Then, Sect. 4 summarizes the COMODO algorithm combining both community detection and subgroup discovery in a

M. Atzmueller (✉)

Tilburg Center for Cognition and Communication (TiCC), Tilburg University, Tilburg, Netherlands

Research Center for Information System Design (ITeG), University of Kassel, Kassel, Germany
e-mail: m.atzmuller@uvt.nl

description-oriented approach [12, 21], for which we also describe an extension for sequential pattern mining. Finally, we conclude with a summary and point out interesting future directions in Sect. 5.

2 Subgroup Discovery

In general, subgroup discovery can be applied for any standard dataset in tabular form in a straight-forward manner using available efficient algorithms, e.g., [6], as implemented in the VIKAMINE [10, 13] platform. Also, for compositional analysis of social networks, i.e., where nodes have attached attribute information, we can directly apply subgroup discovery for identifying interesting subgroups of nodes according to a given quality measure. The description space is then given by all the compositional variables and their respective value domains. As we will see below, it is also possible to combine a structural with a compositional analysis of a network, resulting in description-oriented community detection using subgroup discovery.

2.1 Patterns and Subgroups

Basic concepts used in subgroup discovery [6, 35, 67] are patterns and subgroups. Intuitively, a *pattern* describes a *subgroup*, i.e., the subgroup consists of instances that are covered by the respective pattern. It is easy to see that a pattern describes a fixed set of instances (subgroup), while a subgroup can also be described by different patterns, covering the subgroup's instances. Below, we define these concepts more formally.

A *database* $D = (I, A)$ is given by a set of individuals I and a set of attributes A . A *selector* or *basic pattern* $\text{sel}_{a_i=v_j}$ is a Boolean function $I \rightarrow \{0, 1\}$ that is true if the value of attribute $a_i \in A$ is equal to v_j for the respective individual. For a numeric attribute a_{num} whose range is divided into intervals $e_j = [\min_j, \max_j]$ selectors $\text{sel}_{a_{\text{num}} \in [\min_j; \max_j]}$ can be defined for each interval $[\min_j; \max_j]$ in the domain of a_{num} . The Boolean function is then set to true if the value of attribute a_{num} is within the respective interval. The set of all basic patterns is denoted by S .

Definition 1 A *subgroup description* or (complex) *pattern* sd is given by a set of basic patterns $sd = \{\text{sel}_1, \dots, \text{sel}_l\}$, where $\text{sel}_i \in S$, which is interpreted as a conjunction, i.e., $sd(I) = \text{sel}_1 \wedge \dots \wedge \text{sel}_l$, with $\text{length}(sd) = l$.

Without loss of generality, we focus on a conjunctive pattern language using nominal attribute-value pairs as defined above in this paper; internal disjunctions can also be generated by appropriate attribute-value construction methods, if necessary [14]. We call a pattern p a *superpattern* (or *refinement*) of a *subpattern* p_s , iff $p_s \subset p$.

Definition 2 A subgroup (extension)

$$\text{sg}_{sd} := \text{ext}(sd) := \{i \in I \mid \text{sd}(i) = \text{true}\}$$

is the set of all individuals which are covered by the pattern sd .

As search space for subgroup discovery the set of all possible patterns 2^S is used, that is, all combinations of the basic patterns contained in S . Then, appropriate efficient algorithms, e.g., [6] can be applied.

2.2 Interestingness of a Pattern

A large number of quality functions have been proposed in the literature, see [29] for a comprehensive list, in order to estimate the interestingness of a pattern selected according to the analysis task.

Definition 3 A quality function $q: 2^S \rightarrow \mathbb{R}$ maps every pattern in the search space to a real number that reflects the interestingness of a pattern (or the extension of the pattern, respectively).

Many quality functions for a single target concept (e.g., binary [6, 35] or numerical [6, 43]) trade off the size $n = |\text{ext}(sd)|$ of a subgroup for the deviation $t_{sd} - t_0$, where t_{sd} is the average value of a given target concept in the subgroup identified by the pattern sd and t_0 the average value of the target concept in the general population. In the binary case, the averages relate to the *share* of the target concept. Thus, typical quality functions are of the form

$$q_a(sd) = n^a \cdot (t_{sd} - t_0), \quad a \in [0; 1]. \quad (1)$$

For binary target concepts, this includes, for example, the *weighted relative accuracy* for the size parameter $a = 1$ or a simplified binomial function, for $a = 0.5$. *Multi-target concepts*, e.g., [6, 20, 36, 37] that define a target concept captured by a set of variables can be defined similarly, e.g., by extending a univariate statistical test to the multivariate case, e.g., [20]: Then, the multivariate distributions of a subgroup and the general population are compared in order to identify interesting patterns.

While a quality function provides a *ranking* of the discovered subgroup patterns, often also a statistical assessment of the patterns is useful in data exploration. Quality functions that directly apply a statistical test, for example, the Chi-square quality function, e.g., [6] provide a p -value for simple interpretation. However, the Chi-square quality function estimates deviations in two directions. An alternative, which can also be directly mapped to a p -value is given by the *adjusted residual* quality function q_r , since the values of q_r follow a large standard normal distribution [3]:

$$q_r = n(t_{sd} - t_0) \cdot \frac{1}{\sqrt{nt_0(1-t_0)(1-\frac{n}{N})}} \quad (2)$$

The result of top- k subgroup discovery is the set of the k patterns sd_1, \dots, sd_k , where $sd_i \in 2^S$, with the highest interestingness according to the applied quality function. A subgroup discovery task can now be specified by the five-tuple: (D, c, S, q, k) , where c indicates the target concept; the search space 2^S is defined by the set of basic patterns S .

For several quality functions *optimistic estimates* [6, 31] can be applied for determining upper quality bounds: Consider the search for the k best subgroups: If it can be proven that no subset of the currently investigated hypothesis is interesting enough to be included in the result set of k subgroups, then we can skip the evaluation of any subsets of this hypothesis, but can still guarantee the optimality of the result. More formally, an optimistic estimate $oe(q)$ of a quality function q is a function such that $p \subseteq p' \rightarrow (oe(q))(p) \geq q(p')$, i.e., such that no refinement p' of the pattern p can exceed the quality obtained by $(oe(q))(p)$.

2.3 Community Detection

Communities and cohesive subgroups have been extensively studied in social sciences, e.g., using social network analysis methods [66]. Community detection methods can be classified according to several dimensions, e.g., disjoint vs. overlapping communities. Here, actors in a network can only belong to exactly one community, or to multiple communities at the same time. Furthermore, we distinguish between methods that work on extended (attributed) graphs, i.e., including descriptive information about the nodes. Below, we provide an overview on representative methods, including several basic methods working on simple graphs. After that, we elaborate on methods for detecting overlapping communities, before we focus on descriptive methods.

2.3.1 Basics of Community Detection

Wasserman and Faust [66] discuss social network analysis in depth and provide an overview on the analysis of subgroups/communities in graphs, including clique-based, degree-based, and matrix-perturbation-based methods. Furthermore, several algorithms for community detection have been proposed, formalizing the notions of interesting community structures, and introducing the modularity quality measure [51–53]. Fortunato [26] presents a thorough survey on the state-of-the-art community detection algorithms in graphs, focusing on detecting *disjoint* communities.

For assessing the quality of a community, usually not only the density of the community is assessed but also the connection density of the community is compared to the density of the rest of the network [53]. For the modularity measure the number of connections within the community is compared to the statistically “expected”

number based on all available connections in the network. Besides modularity, prominent examples of community quality measures include, for example, the segregation index [27] and the inverted average out-degree fraction [70].

2.3.2 Detecting Overlapping Communities

Overlapping communities allow an extended modeling of actor–actor relations in social networks: Nodes of a corresponding graph can then participate in multiple communities. This is also typically observed in real-world networks regarding different complementary facets of social interactions [55]. A general overview on algorithms for overlapping community detection is provided by Xie et al. [69]. For example, clique percolation methods proposed in [55, 56] detect k -cliques and then merge them into overlapping communities. Xie and Szymanski [68] present methods that extend the idea of label propagation [58]. Lancichinetti et al. [40] describe an approach for overlapping and hierarchical community structure using a local community metric. The presented metric itself is computed locally but still assesses a global clustering. Further statistical and local optimization algorithms include the COPRA [30] algorithm by Gregory using label propagation of neighboring nodes until a consensus is reached, and the MOSES [46] algorithm by McDaid and Hurley using statistical model-based techniques. Concerning quality measures, extensions of the modularity metric for handling overlapping communities are described in [45, 50, 54].

2.4 Exceptional Model Mining

A general framework for multi-target quality functions in subgroup discovery is given by *exceptional model mining* [6, 41]: It tries to identify interesting patterns with respect to a local model derived from *a set* of attributes. The interestingness can be defined, e.g., by a significant deviation from a model that is derived from the total population or the respective complement set of instances within the population.

In general, a model consists of a specific *model class* and *model parameters* which depend on the values of the model attributes in the instances of the respective pattern cover. The quality measure q then determines the interestingness of a pattern according to its model parameters. Following [42], we outline some simple examples below, focusing on relations between pairs (correlation) and sets of variables (logistic regression):

- A relatively simple example for an exceptionality measure considers the task of identifying subgroups in which the correlation between two numeric attributes is especially strong, e.g., as measured by the Pearson correlation coefficient. This *correlation model class* has exactly one parameter, i.e., the correlation coefficient.

- Furthermore, using a *simple linear regression model*, we can compare the slopes of the regression lines of the subgroup to the general population or the subgroups' complement. This *simple linear regression model* shows the dependency between two numeric variables x and y : It is built by fitting a straight line in the two dimensional space by minimizing the squared residuals e_j of the model:

$$y_i = a + b \cdot x_i + e_j$$

The slope

$$b = \frac{\text{cov}(x, y)}{\text{var}(x)}$$

computed given the covariance $\text{cov}(x, y)$ of x and y , and the variance $\text{var}(x)$ of x can then be used for identifying interesting patterns [41].

- The *logistic regression model* is used for the classification of a binary target attribute $y \in T$ from a set of independent binary attributes $x_j \in T \setminus y, j = 1, \dots, |T| - 1$. The model is given by:

$$y = \frac{1}{1 + e^{-z}}, z = b_0 + \sum_j b_j x_j.$$

Interesting patterns are then those, for example, for which the model parameters b_j differ significantly from those derived from the total population.

Considering network structures, we can also adapt exceptional model mining to that setting. Essentially, it can be regarded as a description-oriented approach for assessing network structures, if the patterns are used to induce graphs or subgraphs. As we will discuss below, we can then also apply exceptional model mining for descriptive community detection, in essence combining subgroup discovery and community detection into a unified approach.

Below, we first outline a quality function for comparing graph structures that correspond to individual patterns (QAP). After that, we discuss quality functions used in community detection in order to assess subgraphs that are induced by some criterion, e.g., by a descriptive pattern.

For some notation, we follow the notions presented in [21]: As outlined above, the concept of a *community* intuitively describes a group C of individuals out of a population such that members of C are strongly “related” among each other but weakly “related” to individuals outside of C . By intuition, this relates, for example, to strongly connected groups of actors in social networks. This idea translates to communities as vertex sets $C \subseteq V$ of a graph $G = (V, E)$. To determine the amount of relatedness (or connectedness, and thus, the community quality of such a subset) several measures have been proposed.

For further concepts regarding our terminology and also the standard community quality functions outlined below, we follow the notation introduced in [21]: For a given undirected graph $G = (V, E)$ and a community $C \subseteq V$: $n := |V|$, let $m := |E|$, $n_C := |C|$, $m_C := |\{\{u, v\} \in E : u, v \in C\}|$ —the number of *intra-edges* of C , and $\bar{m}_C := |\{\{u, v\} \in E : |\{u, v\} \cap C| = 1\}|$ —the number of *inter-edges* of C . Here, it is also convenient to introduce an *inter-degree* for a node $u \in C$ (that depends on the choice of C) by $\bar{d}_C(u) := |\{\{u, v\} \in E : v \notin C\}|$, counting the number of edges between u and nodes outside of C , and $d(u) := |\{\{u, v\} \in E\}|$ is the degree of node u .

There is a wide range of different community evaluation functions $2^V \rightarrow \mathbb{R}$ for estimating the community quality. In the context of this paper, we focus on *maximizing* local quality functions for single communities (which are induced by specific patterns). Therefore, we consider the inverse of a quality measure in those cases, where the measure itself indicates higher quality by lower values.

- Concerning network structures, we can compare adjacency matrices induced by a specific pattern, see [7]. For the assessment we can apply, for example, the quadratic assignment procedure [39] (QAP): it is a standard approach for comparing network structures, e.g., using a graph correlation measure: For comparing two graphs G_1 and G_2 , it estimates the correlation of the respective adjacency matrices M_1 and M_2 and tests that graph level statistic against a QAP null hypothesis [39].

QAP compares the observed graph correlation of (G_1, G_2) to the distribution of the respective resulting correlation scores obtained on repeated random row and column permutations of the adjacency matrix of G_2 . As a result, we obtain a correlation and a statistical significance level according to the randomized distribution scores.

For deriving a quality measure based on QAP and graph correlation, we compare the reference matrix M_N and the matrix M_P for pattern P :

$$q_Q(P) = \text{QAP}(M_N, M_P) = \frac{\text{cov}(M_N, M_P)}{\sqrt{\text{var}(M_N) \cdot \text{var}(M_P)}},$$

where M_N is the transition matrix induced by some reference model (see [7, 24]), and M_P is the transition matrix induced by pattern P , cov indicates the covariance of the matrices, and $\text{var}(M) = \text{cov}(M, M)$ the variance.

For an in-depth description of QAP, we refer to [39]. Furthermore, for the transition matrix, we refer to [23, 24] for more details on the matrix construction step.

- Regarding the quality of a subgraph induced by a pattern, we can adapt the well-known modularity measure to the idea of assessing the induced subgraph captured by a local pattern, i.e., a community pattern (with an associated subgroup description).

In general, the *modularity* MOD [51–53] of a graph clustering with k communities $C_1, \dots, C_k \subseteq V$ focuses on the number of edges *within* a community and compares that with the *expected* such number given a null-model (i.e., a corresponding random graph where the node degrees of G are preserved). It is given by

$$\text{MOD} = \frac{1}{2m} \sum_{u,v \in V} \left(A_{u,v} - \frac{d(u)d(v)}{2m} \right) \delta(C(u), C(v)), \quad (3)$$

where $C(i)$ denotes for $i \in V$ the community to which node i belongs. $\delta(C(u), C(v))$ is the *Kronecker delta* symbol that equals 1 if $C(u) = C(v)$, and 0 otherwise. So, the *modularity* assesses the community quality of a graph partitioning, but can also be adapted to overlapping communities, e.g., [45, 50, 54] for considering the complete graph structure.

For exceptional model mining, however, we need to consider individual patterns. In order to focus on a subgraph induced by a pattern, the *modularity contribution* of a single community C in a *local context* (subgraph induced by the nodes contained in the community C) can then be computed [52, 54] as:

$$\text{MODL}(C) = \frac{1}{2m} \sum_{u,v \in C} \left(A_{u,v} - \frac{d(u)d(v)}{2m} \right),$$

yielding

$$\text{MODL}(C) = \frac{2m_C}{2m} - \sum_{u,v \in C} \frac{d(u)d(v)}{4m^2} = \frac{m_C}{m} - \sum_{u,v \in C} \frac{d(u)d(v)}{4m^2}.$$

- The *segregation index* SIDX [27] is another prominent measure from community detection. It focuses on the local contribution of the pattern, and compares the number of expected inter-edges to the number of observed inter-edges, normalized by the expectation:

$$\text{SIDX}(C) = \frac{E(\bar{m}_C) - \bar{m}_C}{E(\bar{m}_C)} = 1 - \frac{\bar{m}_C n(n-1)}{2m n_C(n-n_C)} \quad (4)$$

- Finally, the *Inverse Average-ODF (out-degree fraction)* IAODF [70] captures the basic intuition of a community regarding the contained vs. the outgoing edges discussed above. As another local measure, IAODF compares the number of *inter-edges* to the number of all edges of a community C , and averages this for the whole community by considering the fraction for each individual node:

$$\text{IAODF}(C) := 1 - \frac{1}{n_C} \sum_{u \in C} \frac{\bar{d}_C(u)}{d(u)} \quad (5)$$

3 Community Detection and Description

While the community detection methods described above only focus on the graph structure, richer graph representations, i.e., *attributed graphs*, enable approaches that specifically exploit the descriptive information of the labels assigned to nodes and/or edges of the graph. Nodes of a network representing users, for example, can be labeled with tags that the respective users utilized in social bookmarking systems, or nodes (denoting actors) can be labeled with properties of the latter. Then, *explicit descriptions* for the characterization of a community can be provided.

Concerning methods that focus on such descriptions in general, an approach for community detection using features identified by frequent pattern mining is presented in [1]; closed frequent patterns are derived and are then used for creating a social network model based on an entropy analysis. However, the network structure itself is not exploited. Similarly, [63] extracts subgraphs with common itemsets. Given a labeled graph, itemset-sharing subgraphs can then be enumerated. However, this approach also does not consider the density of graphs, nor any community measures.

Focusing on methods for generating *explicit descriptions connected with the graph structure*, we distinguish between two types of approaches: first, methods that mainly work on the graph structure but apply descriptive information for restricting the possible sets of communities; second, methods that mine descriptive patterns for obtaining community candidates evaluated using the graph structure. As a representative of the first type, the concepts of dense subgraphs and subspace clusters for mining cohesive patterns are combined in [49].

Starting with quasi-cliques, these are expanded until constraints regarding the description or the graph structure are violated. Similarly, [32] combines subspace clustering and dense subgraph mining, also interleaving quasi-clique and subspace construction. As an example for the second type outlined above, [28] proposes an approach for the problem of finding overlapping communities in graphs and social networks that aims at detecting the top- k communities such that the total edge density over all k communities is maximized. The three algorithmic variants proposed in [28] apply a greedy strategy for detecting dense subgroups, and restrict the result set of communities, such that each edge can belong to at most one community. This partitioning involves a global approach on the community quality. Furthermore, [64] study the correlation between attribute sets and the occurrence of dense subgraphs in large attributed graphs. The proposed method considers frequent attribute sets using an adapted frequent item mining technique, and identifies the top- k dense subgraphs induced by a particular attribute set, called structural correlation patterns. The DCM method presented in [57] includes a two-step process of community detection and community description. A heuristic approach is applied for discovering the top- k communities. Pool et al. utilize a special interestingness function which is based on counting outgoing edges of a community similar to the IAODF measure; for that, they also demonstrate the trend of a correlation with the modularity function.

Furthermore, the COMODO algorithm [21] that we summarize in the next section combines community detection and subgroup discovery resulting in a description-oriented approach. By specifying a standard quality function the quality of the communities to discover can be estimated. Then, this quality function can be specifically selected according to the analysis task.

4 Community Detection Using Exceptional Model Mining

For providing both structurally valid and interpretable communities we utilize the graph structure as well as additional descriptive features of the nodes. Hence, we identify communities as sets of nodes together with a *description* composed of the nodes' features. Such a *community pattern* then provides an intuitive description of the community, e.g., by an easily interpretable conjunction of attribute-value pairs. Basically, we aim at identifying communities according to standard community quality measures. Below, we first provide an algorithmic overview on the approach and summarize exemplary evaluation results. After that, we sketch the application of the algorithm for community detection on dynamic networks, i.e., for identifying exceptional sequential patterns.

4.1 COMODO: *Description-Oriented Community Detection*

Below, we summarize the COMODO algorithm presented in [21]: It focuses on *description-oriented community detection* using subgroup discovery, and aims at discovering the top- k communities (described by community patterns). The method is based on an adapted subgroup discovery approach [12, 42], and also tackles typical problems that are not addressed by standard approaches for community detection such as pathological cases like small community sizes. COMODO utilizes optimistic estimates [31, 67], which are efficient to compute, in order to prune the search space significantly. For that, a number of standard community evaluation functions have been applied using optimistic estimates for an efficient approach.

4.1.1 Algorithmic Overview

COMODO utilizes both the graph structure and descriptive information of the attributed graph. This information is contained in two data structures: The graph structure is encoded in graph G while the attribute information is contained in database D describing the respective attribute values of each node. In a preprocessing step, we merge these data sources. Since the communities considered in our approach do not contain isolated nodes, we can describe them as sets of edges. We transform the data (of the given graph G and the database D containing the nodes' descriptive information) into a new dataset focusing on the edges of the graph G :

Each data record in the new dataset represents an edge between two nodes. The attribute values of each such data record are the common attributes of the edge’s two nodes. For a more detailed description, we refer to [21].

COMODO utilizes an extended FP-tree (frequent pattern tree) structure inspired by the FP-growth algorithm, which compiles the data in a convenient prefix pattern tree structure for mining frequent itemsets, see [2] for a detailed description. Our adapted tree structure is called the *community pattern tree* (CP-tree) that allows to efficiently traverse the solution space. The tree is built in two scans of the graph dataset and is then mined in a recursive divide-and-conquer manner, see [9, 42] for more details. In the main algorithmic procedure of COMODO, patterns containing only one basic pattern are mined first. Then, patterns conditioned on the occurrence of a (prefixed) complex pattern (as a set of basic patterns, chosen in the previous recursion step) are considered recursively. For more algorithmic details, we refer to [21]. As described there, we can apply standard quality functions efficiently using optimistic estimates, e.g., for the *modularity* or the *segregation index*, see [21] for more details.

4.1.2 Illustrative Evaluation Results

Below, we present illustrative evaluation results [21] considering the efficiency of the applied optimistic estimates, and the validity of the obtained patterns. For that, we compared the total number of search steps, that is community allocations that are considered by the COMODO algorithm, with no optimistic estimate pruning to optimistic estimate pruning using different community quality measures. Additionally, we measured the impact of using different minimal community size thresholds. Some results are shown in Fig. 1 for the BibSonomy click graph for

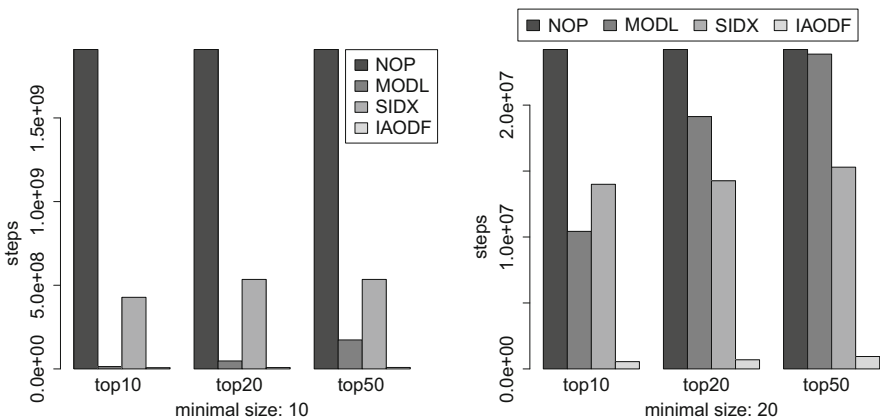


Fig. 1 Runtime performance of COMODO on the BibSonomy click graph, see [21] for more details: search steps with no optimistic estimate pruning (*NOP*) vs. community quality functions with optimistic estimate pruning: *MODL* (Local modularity), *SIDX* (Segregation Index), and *IAODF* (Inverse Average-ODF), for minimal size thresholds $\tau_n = 10, 20$

$k = 10, 20, 50$ and minimal size thresholds $\tau_n = 10, 20$. We consider a number of standard community quality functions, that is, the *segregation index*, the *Inverse Average-ODF*, and the *modularity*.

The large, exponential search space can be exemplified, e.g., for the click graph with a total of about $2 \cdot 10^{10}$ search steps for a minimal community size threshold $\tau_n = 10$. The results demonstrate the effectiveness of the proposed descriptive mining approach applying the presented optimistic estimates. The implemented pruning scheme makes the approach scalable for larger datasets, especially when the local modularity quality function is chosen to assess the communities' quality. Concerning the validity of the patterns, we focused on structural properties of the patterns and the subgraphs induced by the respective community patterns. We applied the significance test described in [38] for testing the statistical significance of the density of a discovered subgraph. Furthermore, we compared COMODO to three baseline community detection algorithms [30, 46, 57], where COMODO consistently shows a significantly better performance concerning validity and description length; for more details, we refer to [21].

4.2 Sequential Pattern Analysis: Detecting Exceptional Link Trails

In addition to static community detection, we can also consider temporal aspects, i.e., focusing on sequences of states or events which can be applied for a variety of analysis ranging from the analysis of human behavior [23] to industrial applications [24]. In an extended modeling approach, we can map transitions between states to a weighted network, according to a first-order Markov chain model. Below, we outline an approach for detecting exceptional sequential link trails captured by community patterns, see [7] for a detailed description.

As before, our subject of analysis is given by an attributed graph that models the link trails in the following way: Nodes of the graph denote actors of a social network, e.g., users of a social system or locations in a location-based social network. The edges of the graph model the links between the nodes (as transitions). As a simple example, we can consider a set of users and a set of locations. Each user visits a sequence of locations—in a location-based social network. Then, we are interested in modeling these sequences (of locations), and in detecting exceptional groups of transitions (between locations) w.r.t. users and their properties, respectively.

At a music event festival, for example, possible characterizing factors describing certain users groups could be specific music genres. Here, exceptional patterns could include, for example, users being interested in *rock music* and *dance* visiting only a very specific selection of performances in characteristic sequences, compared to the behavior of all users and their sequential link trails. Essentially, we apply descriptive community detection (e.g., using COMODO) on the attributed graph, where the edges indicate transitions between states according to a first-order Markov chain modeling approach [44, 65].

4.2.1 Modeling

For our attributed graph model, we label the links according to the descriptive information of the sequential trail. Then, we identify exceptional community patterns based on the labels and structure of the contained links using exceptional model mining. In particular, we assess a pattern capturing a set of nodes that model the state space of the respective transitions.

For constructing a reference model, we construct transition matrices corresponding to the *observed data*. For those observed sequences we can simply construct transition matrices counting the transitions between the individual states. We construct an according matrix M^N with $m_{ij}^N = |\text{suc}(i, j)|$, where $\text{suc}(i, j)$ denotes the successive sequences from state i to state j contained in the sequence.

A community pattern P induces a subgraph (community) C_P given a set of labels P , selecting all links that are covered, i.e., that share a label contained in P . Then, all transitions in the matrix M^N are selected (corresponding to a set of links of the network) that are covered by the pattern P . Using that, we construct an according transition pattern matrix M^P based on the respective counts of the covered transitions. Intuitively, the matrix M^P can then be regarded as some kind of “projection” of matrix M^N given the pattern P using our modeling approach. In the simplest case, we can just transfer the weighted links of the subgraph C_P . For identifying exceptional models (M_P induced by P) we can then apply, e.g., the QAP quality function $q_Q(P) = \text{QAP}(M_N, M_P)$ introduced above.

4.2.2 Results

For some illustrative results (see [7] for more details), we utilized data from the EveryAware¹ project, e.g., [19]. Specifically, we focused on collectively organized noise measurements collected using the *WideNoise Plus* application between December 14, 2011 and June 6, 2014, see [20] for more details. *WideNoise Plus* allows the collection of noise measurements using smartphones. It includes sensor data from the microphone given as noise level in dB(A), the location from the GPS-, GSM-, and WLAN-sensor represented as latitude and longitude coordinate, as well as a timestamp. In addition, tags can be assigned to the recording. We collected data from all around the world using iOS and Android devices.

In total, the applied dataset contains 6069 data records, i.e., noise measurements of 635 users (i.e., 635 trails, with an average trail length of about 10) and 2009 distinct tags. Table 1 shows exemplary exceptional conforming and deviating patterns using q_Q as quality measure. In addition, it shows the sizes of the covered subsets. From a qualitative point of view, the patterns shown in the table are intuitive

¹<http://www.everyaware.eu>.

Table 1 Illustrative exceptional conforming/deviating community patterns for *WideNoise Plus*

#	q_Q	Size	Description
1	0.94	5078	Traffic
2	0.89	3990	Car
3	0.76	3326	Noise
4	0.43	707	Bird \wedge courtyard
5	0.24	600	Background \wedge quiet

Patterns #1–#3 tend rather to conform to the reference model (especially #1 and #2), while patterns #4–#5 (increasingly) show a deviating behavior

to interpret and also tend to conform to our expectations concerning the reference behavior of the dataset, where we can clearly identify deviations concerning noisy and relatively quiet environments.

5 Conclusions

In this paper, we have presented an organized view on descriptive community detection. Specifically, we described subgroup discovery for compositional network analysis concerning properties of the actors, with extensions to the analysis of complex target concepts like correlations between a set of variables, or dense subgraphs—captured by exceptional model mining approaches. Then, this directly extends to community detection on attributed graphs. In particular, we summarized the COMODO algorithm that combines community detection and exceptional model mining, resulting in a description-oriented approach for community analytics. We furthermore sketched an extension to dynamic data, considering sequential patterns capturing exceptional sequential link trails. This adds one further dimension to the descriptive approaches, by considering by static as well as dynamic phenomena, and enables the modeling and investigation of complex analysis tasks.

For future work, we aim to extend the analysis towards further time-oriented representations, e.g., considering sequences of graphs, and the evolution of communities, e.g., [33, 34]. Also, we aim to integrate and exploit methods for generating descriptions and the respective relations in link analytics, e.g., in link prediction [60–62] on multiplex networks. Then, besides the detection of communities, also their analysis and assessment in the form of descriptive patterns is highly relevant, e.g., [11, 15, 17, 18] also concerning their semantic grounding [47, 48], and integration into explanation-aware approaches [16, 25, 59]. Furthermore, developing scalable methods for enabling such approaches for large and complex datasets, e.g., [22, 42] is another interesting direction for future work.

References

1. Adnan, M., Alhaji, R., Rokne, J.: Identifying social communities by frequent pattern mining. In: Proc. 13th Intl. Conf. Information Visualisation, pp. 413–418. IEEE Computer Society, Washington, DC (2009)
2. Agrawal, R., Srikant, R.: Fast algorithms for mining association rules. In: Bocca, J.B., Jarke, M., Zaniolo, C. (eds.) Proc. 20th Int. Conf. Very Large Data Bases (VLDB), pp. 487–499. Morgan Kaufmann, San Francisco (1994)
3. Agresti, A.: An Introduction to Categorical Data Analysis. Wiley, Hoboken (2007)
4. Atzmueller, M.: Data mining on social interaction networks. *J. Data Min. Digit. Humanit.* **1**, pp. 1–34 (2014)
5. Atzmueller, M.: Subgroup and community analytics on attributed graphs. In: Kuznetsov, S.O., Missaoui, R., Obiedkov, S. (eds.) Proceedings of the International Workshop on Social Network Analysis Using Formal Concept Analysis (SNAFCA-2015), CEUR-WS, vol. 1534 (2015)
6. Atzmueller, M.: Subgroup discovery – advanced review. *WIREs Data Min. Knowl. Discov.* **5**(1), 35–49 (2015)
7. Atzmueller, M.: Detecting community patterns capturing exceptional link trails. In: Proceedings of the IEEE/ACM ASONAM. IEEE Press, Boston, MA (2016)
8. Atzmueller, M.: Local exceptionality detection on social interaction networks. In: Proceedings of the ECML-PKDD 2016: European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Springer, Berlin (2016)
9. Atzmueller, M., Lemmerich, F.: Fast subgroup discovery for continuous target concepts. In: Proceedings of the International Symposium on Methodologies for Intelligent Systems. LNCS, vol. 5722, pp. 1–15. Springer, Heidelberg (2009)
10. Atzmueller, M., Lemmerich, F.: VIKAMINE - open-source subgroup discovery, pattern mining, and analytics. In: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Springer, Heidelberg (2012)
11. Atzmueller, M., Lemmerich, F.: Exploratory pattern mining on social media using geo-references and social tagging information. *Int. J. Web Sci.* **2**(1/2), 80–112 (2013)
12. Atzmueller, M., Mitzlaff, F.: Efficient descriptive community mining. In: Proceedings of the 24th International FLAIRS Conference, pp. 459–464. AAAI Press, Palo Alto, CA (2011)
13. Atzmueller, M., Puppe, F.: Semi-automatic visual subgroup mining using VIKAMINE. *J. Univers. Comput. Sci.* **11**(11), 1752–1765 (2005)
14. Atzmueller, M., Puppe, F.: SD-Map - a fast algorithm for exhaustive subgroup discovery. In: Proceedings of the European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD), pp. 6–17. Springer, Heidelberg (2006)
15. Atzmueller, M., Puppe, F.: A case-based approach for characterization and analysis of subgroup patterns. *J. Appl. Intell.* **28**(3), 210–221 (2008)
16. Atzmueller, M., Roth-Berghofer, T.: The mining and analysis continuum of explaining uncovered. In: Proceedings of the 30th SGAI International Conference on Artificial Intelligence (AI-2010) (2010)
17. Atzmueller, M., Baumeister, J., Hemsing, A., Richter, E.J., Puppe, F.: Subgroup mining for interactive knowledge refinement. In: Proceedings of the 10th Conference on Artificial Intelligence in Medicine (AIME 05). LNAI, vol. 3581, pp. 453–462. Springer, Heidelberg (2005)
18. Atzmueller, M., Baumeister, J., Puppe, F.: Introspective subgroup analysis for interactive knowledge refinement. In: Proceedings of the 19th International Florida Artificial Intelligence Research Society Conference 2006 (FLAIRS-2006), pp. 402–407. AAAI Press, Palo Alto, CA (2006)

19. Atzmueller, M., Becker, M., Kibanov, M., Scholz, C., Doerfel, S., Hotho, A., Macek, B.E., Mitzlaff, F., Mueller, J., Stumme, G.: Ubicon and its applications for ubiquitous social computing. *New Rev. Hypermedia Multimed.* **20**(1), 53–77 (2014)
20. Atzmueller, M., Mueller, J., Becker, M.: Exploratory subgroup analytics on ubiquitous data. In: *Mining, Modeling and Recommending ‘Things’ in Social Media*. LNAI, vol. 8940. Springer, Heidelberg (2015)
21. Atzmueller, M., Doerfel, S., Mitzlaff, F.: Description-oriented community detection using exhaustive subgroup discovery. *Inf. Sci.* **329**, 965–984 (2016)
22. Atzmueller, M., Mollenhauer, D., Schmidt, A.: Big data analytics using local exceptionality detection. In: *Enterprise Big Data Engineering, Analytics, and Management*. IGI Global, Hershey, PA (2016)
23. Atzmueller, M., Schmidt, A., Kibanov, M.: DASHTrails: an approach for modeling and analysis of distribution-adapted sequential hypotheses and trails. In: *Proceedings of the WWW 2016 (Companion)*, IW3C2/ACM (2016)
24. Atzmueller, M., Schmidt, A., Kloeppe, B., Arnu, D.: HypGraphs: an approach for modeling and comparing graph-based and sequential hypotheses. In: *Proceedings of the ECML-PKDD Workshop on New Frontiers in Mining Complex Patterns (NFMCP)*, Riva del Garda (2016)
25. Clancey, W.J.: The epistemology of a rule-based expert system: a framework for explanation. *Artif. Intell.* **20**, 215–251 (1983)
26. Fortunato, S.: Community detection in graphs. *Phys. Rep.* **486**(3–5), 75–174 (2010)
27. Freeman, L.: Segregation in social networks. *Sociol. Methods Res.* **6**(4), 411 (1978)
28. Galbrun, E., Gionis, A., Tatti, N.: Overlapping community detection in labeled graphs. *Data Min. Knowl. Discov.* **28**(5–6), 1586–1610 (2014)
29. Geng, L., Hamilton, H.J.: Interestingness measures for data mining: a survey. *ACM Comput. Surv.* **38**(3), Article No. 9 (2006)
30. Gregory, S.: Finding overlapping communities in networks by label propagation. *New J. Phys.* **12**, 103018 (2010)
31. Grosskreutz, H., Rüping, S., Wrobel, S.: Tight optimistic estimates for fast subgroup discovery. In: *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. LNCS, vol. 5211, pp. 440–456. Springer, Heidelberg (2008)
32. Günnemann, S., Färber, I., Boden, B., Seidl, T.: GAMer: a synthesis of subspace clustering and dense subgraph mining. In: *Knowledge and Information Systems*. Springer, London (2013)
33. Kibanov, M., Atzmueller, M., Scholz, C., Stumme, G.: Temporal evolution of contacts and communities in networks of face-to-face human interactions. *Sci. China* **57**, 1–17 (2014)
34. Kibanov, M., Atzmueller, M., Illig, J., Scholz, C., Barrat, A., Cattuto, C., Stumme, G.: Is web content a good proxy for real-life interaction? A case study considering online and offline interactions of computer scientists. In: *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE Press, Boston, MA (2015)
35. Klösgen, W.: Explora: a multipattern and multistrategy discovery assistant. In: Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P., Uthurusamy, R. (eds.) *Advances in Knowledge Discovery and Data Mining*, pp. 249–271. AAAI Press, Menlo Park (1996)
36. Klösgen, W.: 16.3: subgroup discovery. In: *Handbook of Data Mining and Knowledge Discovery*. Oxford University Press, New York (2002)
37. Klösgen, W.: 5.2: subgroup patterns. In: *Handbook of Data Mining and Knowledge Discovery*. Oxford University Press, New York (2002)
38. Koyuturk, M., Szpankowski, W., Grama, A.: Assessing significance of connectivity and conservation in protein interaction networks. *J. Comput. Biol.* **14**(6), 747–764 (2007)
39. Krackhardt, D.: QAP partialling as a test of spuriousness. *Soc. Netw.* **9**, 171–186 (1987)
40. Lancichinetti, A., Fortunato, S., Kertsz, J.: Detecting the overlapping and hierarchical community structure in complex networks. *New J. Phys.* **11**(3), 033015 (2009)

41. Leman, D., Feelders, A., Knobbe, A.: Exceptional model mining. In: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Lecture Notes in Computer Science, vol. 5212, pp. 1–16. Springer, Heidelberg (2008)
42. Lemmerich, F., Becker, M., Atzmueller, M.: Generic pattern trees for exhaustive exceptional model mining. In: Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases. Springer, Heidelberg (2012)
43. Lemmerich, F., Atzmueller, M., Puppe, F.: Fast exhaustive subgroup discovery with numerical target concepts. *Data Min. Knowl. Discov.* **30**, 711–762 (2016)
44. Lempel, R., Moran, S.: The stochastic approach for link-structure analysis (SALSA) and the TKC effect. *Comput. Netw.* **33**(1), 387–401 (2000)
45. Lin, Y.R., Chi, Y., Zhu, S., Sundaram, H., Tseng, B.L.: Analyzing communities and their evolutions in dynamic social networks. *ACM Trans. Knowl. Discov. Data* **3**, 8:1–8:31 (2009)
46. McDaid, A., Hurlley, N.: Detecting highly overlapping communities with model-based overlapping seed expansion. In: Proceedings of the International Conference on Advances in Social Networks Analysis and Mining, ASONAM, pp. 112–119. IEEE Computer Society, Washington, DC (2010)
47. Mitzlaff, F., Atzmueller, M., Stumme, G., Hotho, A.: Semantics of user interaction in social media. In: Ghoshal, G., Poncela-Casasnovas, J., Tolksdorf, R. (eds.) *Complex Networks IV. Studies in Computational Intelligence*, vol. 476. Springer, Heidelberg (2013)
48. Mitzlaff, F., Atzmueller, M., Hotho, A., Stumme, G.: The social distributional hypothesis. *J. Soc. Netw. Anal. Min.* **4**, 216 (2014)
49. Moser, F., Colak, R., Rafiey, A., Ester, M.: Mining cohesive patterns from graphs with feature vectors. In: *SDM, SIAM*, vol. 9, pp. 593–604 (2009)
50. Muff, S., Rao, F., Cafilisch, A.: Local modularity measure for network clusterizations. *Phys. Rev. E Stat. Nonlinear Matter Phys.* **72**(5), 056107 (2005)
51. Newman, M.E.J.: Detecting community structure in networks. *Eur. Phys. J.* **38**, 321–330 (2004)
52. Newman, M.E.J.: Modularity and community structure in networks. *Proc. Natl. Acad. Sci.* **103**(23), 8577–8582 (2006)
53. Newman, M.E.J., Girvan, M.: Finding and evaluating community structure in networks. *Phys. Rev. E Stat. Nonlin Soft Matter Phys.* **69**(2), 1–15 (2004)
54. Nicosia, V., Mangioni, G., Carchiolo, V., Malgeri, M.: Extending the definition of modularity to directed graphs with overlapping communities. *J. Stat. Mech.* **2009**, 03024 (2009)
55. Palla, G., Dernyi, I., Farkas, I., Vicsek, T.: Uncovering the overlapping community structure of complex networks in nature and society. *Nature* **435**(7043), 814–818 (2005)
56. Palla, G., Farkas, I.J., Pollner, P., Derenyi, I., Vicsek, T.: Directed network modules. *New J. Phys.* **9**(6), 186 (2007)
57. Pool, S., Bonchi, F., van Leeuwen, M.: Description-driven community detection. *Trans. Intell. Syst. Technol.* **5**(2), 1–21 (2014)
58. Raghavan, U., Albert, R., Kumara, S.: Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* **76**, 036106 (2007)
59. Roth-Berghofer, T.R., Cassens, J.: Mapping goals and kinds of explanations to the knowledge containers of case-based reasoning systems. In: Muñoz-Avila, H., Ricci, F. (eds.) *Case-Based Reasoning Research and Development*, 6th International Conference on Case-Based Reasoning, ICCBR 2005, Chicago, IL, USA, August 2005, Proceedings. Lecture Notes in Artificial Intelligence, vol. 3620, pp. 451–464. Springer, Heidelberg (2005)
60. Scholz, C., Atzmueller, M., Stumme, G.: On the predictability of human contacts: influence factors and the strength of stronger ties. In: Proceedings of the 4th ASE/IEEE International Conference on Social Computing (SocialCom). IEEE Computer Society, Boston, MA (2012)
61. Scholz, C., Atzmueller, M., Barrat, A., Cattuto, C., Stumme, G.: New Insights and Methods For Predicting Face-To-Face Contacts. In: Kiciman E, Ellison NB, Hogan B, Resnick P, Soboroff I (eds) *Proc. International AAAI Conference on Weblogs and Social Media*. AAAI Press, Palo Alto, CA (2013)

62. Scholz, C., Atzmueller, M., Kibanov, M., Stumme, G.: How do people link? Analysis of contact structures in human face-to-face proximity networks. In: Proc. ASONAM 2013. ACM Press, New York, NY (2013)
63. Sese, J., Seki, M., Fukuzaki, M.: Mining networks with shared items. In: Proceedings of the 19th ACM International Conference on Information and Knowledge Management, pp. 1681–1684. ACM, New York, NY (2010)
64. Silva, A., Meira Jr., W., Zaki, M.J.: Mining attribute-structure correlated patterns in large attributed graphs. Proc VLDB Endowment **5**(5), 466–477 (2012)
65. Singer, P., Helic, D., Taraghi, B., Strohmaier, M.: Detecting memory and structure in human navigation patterns using Markov chain models of varying order. PLoS One **9**(7), e102070 (2014)
66. Wasserman, S., Faust, K.: Social Network Analysis: Methods and Applications. Structural Analysis in the Social Sciences, vol. 8, 1st edn. Cambridge University Press, Cambridge (1994)
67. Wrobel, S.: An algorithm for multi-relational discovery of subgroups. In: Proceedings of the 1st European Symposium on Principles of Data Mining and Knowledge Discovery, pp. 78–87. Springer, Heidelberg (1997)
68. Xie, J., Szymanski, B.K.: LabelRank: a stabilized label propagation algorithm for community detection in networks. In: Proceedings of the IEEE Network Science Workshop, West Point, NY (2013)
69. Xie, J., Kelley, S., Szymanski, B.K.: Overlapping community detection in networks: the state-of-the-art and comparative study. ACM Comput. Surv. **45**(4), 43:1–43:35 (2013)
70. Yang, J., Leskovec, J.: Defining and evaluating network communities based on ground-truth. In: Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics, MDS '12, pp. 3:1–3:8. ACM, New York, NY (2012)