

# 1 Predicting the Evolution of Communities with 2 Online Inductive Logic Programming

3 **George Athanasopoulos**<sup>1</sup>  
4 geotha1995@hotmail.com

5 **George Paliouras**<sup>2</sup>  
6 paliourg@iit.demokritos.gr

7 **Dimitrios Vogiatzis**<sup>3,2</sup>  
8 dimitrv@iit.demokritos.gr

9 **Grigorios Tzortzis**<sup>2</sup>  
10 gtzortzi@iit.demokritos.gr

11 **Nikos Katzouris**<sup>2</sup>  
12 nkatz@iit.demokritos.gr

13  
14 <sup>1</sup>Department of Informatics and Telecommunications, National and Kapodistrian University of  
15 Athens, Athens, Greece

16 <sup>2</sup>Institute of Informatics and Telecommunications, NCSR “Demokritos”, Athens, Greece

17 <sup>3</sup>The American College of Greece, Deree, Athens, Greece

## 18 — Abstract —

19 In the recent years research on dynamic social network has increased, which is also due to the  
20 availability of data sets from streaming media. Modeling a network’s dynamic behaviour can be  
21 performed at the level of communities, which represent their mesoscale structure. Communities  
22 arise as a result of user to user interaction. In the current work we aim to predict the evolution of  
23 communities, i.e. to predict their future form. While this problem has been studied in the past as  
24 a supervised learning problem with a variety of classifiers, the problem is that the “knowledge” of  
25 a classifier is opaque and consequently incomprehensible to a human. Thus we have employed first  
26 order logic, and in particular the event calculus to represent the communities and their evolution.  
27 We addressed the problem of predicting the evolution as an online Inductive Logic Programming  
28 problem (ILP), where the issue is to learn first order logical clauses that associate evolutionary  
29 events, and particular Growth, Shrinkage, Continuation and Dissolution to lower level events.  
30 The lower level events are features that represent the structural and temporal characteristics of  
31 communities. Experiments have been performed on a real life data set form the Mathematics  
32 StackExchange forum, with the OLED framework for ILP. In doing so we have produced clauses  
33 that model both short term and long term correlations.

34 **2012 ACM Subject Classification** Human-centered computing → Social network analysis, Com-  
35 puting methodologies → Machine learning, Computing methodologies → Online learning settings,  
36 Computing methodologies → Inductive logic learning

37 **Keywords and phrases** Social Network Analysis, Community Evolution Prediction, Machine  
38 Learning, Inductive Logic Programming, Event Calculus, Online Learning

39 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

## 40 **1** Introduction

41 A social network is a structure which contains individuals, who are linked to other individuals.  
42 The link among them states an interaction which has one or more types of interdependency



© George Athanasopoulos, George Paliouras, Dimitrios Vogiatzis, Grigorios Tzortzis, Nikos Katzouris;  
licensed under Creative Commons License CC-BY

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:21

Leibniz International Proceedings in Informatics



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

43 such as friendship, kinship, common interest, financial exchange. Social networks are often  
44 represented as graphs, with nodes representing users and edge representing interactions.  
45 Usually a social network changes over time because new individuals join the network, new  
46 interactions are developed, or some individuals cease to be active for a short or a long period.  
47 This is actually the predominant behaviour especially in streaming social media, such as  
48 forums.

49 The social networks are often studied at the level of communities, which represent their  
50 meso-scale structure. A group of nodes forms a community if it densely connected, and  
51 sparsely connected to other communities. The said communities are not explicitly formed  
52 but rather implicitly as a result of the actions of individual users, that are not random but  
53 tend to follow a certain pattern that is related to their similarity to other users. There are  
54 many algorithms that have been developed for the detection of communities in networks that  
55 are static [6].

56 In dynamic networks, the communities are influenced over the time by its users' interaction.  
57 This influence causes changes in the structure of the communities. Many researchers, consider  
58 that the structure of a community contains important information for network evolution as a  
59 whole. Thus, it is highly imperative to model the dynamic behavior in social networks and  
60 try to predict their evolution.

61 In this paper we study the problem of community evolution prediction in dynamic social  
62 networks. We address this problem as a supervised learning task where we predict four types  
63 of community evolutionary events, *growth*, *shrinkage*, *continuation* and *dissolution*. Various  
64 features were investigated in order to understand how they influence the results. Among  
65 them, are the structural and temporal characteristics of communities. What is unique in  
66 the current approach is that we use a first order logic formalism to represent the correlation  
67 between evolutionary events and the input features. Moreover Inductive Logic Programming  
68 (ILP) us used to learn event calculus clauses. Event calculus was chosen because it is human  
69 understandable, it can be used to model effect of actions in time, and the variation we have  
70 adopted can perform ILP in an online fashion which is especially useful in streaming media.

71 The rest of the paper is organised as follows. In Section 2 we refer to past work on  
72 community evolution prediction, in Section 3 we refer to the Event Calculus as a logic  
73 formalism but also to Inductive Logic Programming as a way of learning clauses, then in  
74 Section 4 we refer to the methodology we followed for community evolution prediction, in  
75 Section 5 we present experimental results, conclusions are drawn in Section 6. In appendix A  
76 we present samples of Event Calculus clauses, and we elaborate on the Hoeffding bound.

## 77 **2 Related Work**

78 One of the most interesting problems in social network prediction is the prediction of  
79 possible events, at the level of communities, such as growing, shrinking, merging with  
80 another community or splitting into more communities The research in community evolution  
81 prediction and the events have been proposed is quite extensive.

82 Patil et al. [14] predicted whether a community will disappear or will survive. They  
83 observed that both the level of member diversity and social activities are critical in maintaining  
84 the stability of communities. They also found that certain prolific members play an important  
85 role in maintaining the community's stability. Goldberg et al. [8] correlated the lifespan  
86 of a community with the structural parameters of its early stages. Brodka et al. [7], [2]  
87 tried to predict 6 evolutionary events of communities, i.e. growth, shrinkage, continuation,  
88 dissolution, merging and split. The used as features the history of the events of a community

89 in the three preceding timeframes, and the community size in these timeframes. They  
 90 found that the prediction based on simple input features may be very accurate, while some  
 91 classifiers are more precise than the others. Kairam et al. [11] tried to understand the factors  
 92 contributing to the growth and longevity in a social network. They investigated the role  
 93 that two types of growth (Diffusion and non-diffusion) play during a community's formative  
 94 stages. Diffusion growth is when a community attracts new members through ties to existing  
 95 members. Non-diffusion growth occurs with individuals with no prior ties to the network.  
 96 Diakidis et al. [4] studied on-line social networks as a supervised learning task with sequential  
 97 and non-sequential classifiers. Structural, content and contextual features as well as the  
 98 previous states of a community are considered as the features that are involved in the task of  
 99 community evolution. The evolution phenomena they tried to predict are the continuation,  
 100 shrinking, growth and dissolution.

101 Takaffoli et al. [16] quantified the events that may occur in a community as follows:  
 102 survive:{true, false}, merge:{true, false}, split:{true, false}, size:{expand, shrink}, and  
 103 cohesion:{cohesive, loose}. First they tried to predict whether a community will survive,  
 104 followed by a separate predictor for each of the events.

105 Ilhan et al. [10] proposed a regression ARIMA model to predict values of community  
 106 features based on the values of the past community instances. Then the predicted community  
 107 features are used to train a classifier to predict the evolutionary events.

### 108 **3 Background: Event Calculus and Inductive logic programming**

109 The Event Calculus (EC) is a temporal logic formalism for reasoning about actions and  
 110 changes. [13] EC that has been used as a basis in event recognition applications, providing  
 111 among others, direct connections to machine learning, via Inductive Logic Programming  
 112 (ILP) [3]. Its ontology comprises *time points*, represented by integers; time varying properties  
 113 known as *fluents*; and actions known as *events*. The events occur in time and may affect the  
 114 fluents by altering their value. The axioms of the EC incorporate the *law of inertia*, according  
 115 to which fluents persist over time, unless they are affected by an event. Thus, if an event  
 116 is initiated at time  $T$ , it will persist until another event will fire a termination rule. Also,  
 117 if an event is terminated at time  $T$ , it will remain terminated until another event fires an  
 118 initiation rule. The basic predicates are presented in Table 1, while the domain-independent  
 119 axioms are in Table 2. Axiom (1) states that a high level event, represented as fluent  $F$  for  
 120 convenience, is happening at time  $T$  if it has been initiated at the previous time point. While  
 121 Axiom (2) states that  $F$  continues to happen unless it is terminated.

122 Let us examine how the evolution of a community could be represented in terms of EC;  
 123 an evolutionary phenomenon such as the *growth* will be represented as a fluent, whereas the  
 124 factors that contribute to the *growth* will be represented as events. For example in Table 3, it  
 125 means that community  $X_0$  will grow in time  $T$  provided its size was 3 and its density was 4.  
 126 Likewise, rules can be formed for the rest of the evolutionary phenomena as combinations of  
 127 features (or events). The rules are also known as *clauses*, whereas the predicates *happensAt*  
 128 are also known as *literals*. In addition, a rule  $B$  is a specialisation of rule  $A$  if the instances  
 129 that satisfy rule  $B$  are a subset of the instances that satisfy rule  $A$ .

130 The problem that we try to address is to learn such rules (or clauses) from data, rather  
 131 than hand encode them; for that we can use inductive logic programming (ILP). In ILP  
 132 programming first order rules are learnt from relational data under a supervised learning  
 133 scheme. Thus we have input data and class labels. The input data are often named the  
 134 *narrative*, and the class label is known as the *annotation*.

■ **Table 1** The basic predicates of the EC

Predicate	Meaning
$\text{happensAt}(E,T)$	Event $E$ occurs at time $T$
$\text{initiatedAt}(F,T)$	At time $T$ fluent $F$ is initiated
$\text{terminatedAt}(F,T)$	At time $T$ fluent $F$ is terminated
$\text{holdsAt}(F,T)$	Fluent $F$ holds at time $T$

■ **Table 2** The domain-independent axioms

Axioms
$\text{holdsAt}(F,T+1) \leftarrow \text{initiatedAt}(F,T). \quad (1)$
$\text{holdsAt}(F,T+1) \leftarrow \text{holdsAt}(F,T), \text{not terminatedAt}(F,T). \quad (2)$

135 Given an encoding of the known background knowledge and a set of examples represented  
 136 as a logical database of facts, an ILP system will derive a *hypothesised* logic program which  
 137 entails all the positive and none of the negative examples. ILP provides various techniques  
 138 for learning logical theories from examples. In Learning from Interpretations (LFI) [1] setting  
 139 each training example is an interpretation, i.e. a set of narrative and annotation atoms (see  
 140 Table 4). Given a set of training interpretations  $I$  and some background theory  $B$ , which  
 141 consists of the domain-independent axioms of the EC, the goal in LFI is to find a theory.

142 In this paper, OLED (Online Learning of Event Definitions) [12] was used for learning rules  
 143 that perform community evolution prediction. OLED is an online ILP system for learning  
 144 logical theories from data streams. It has been designed having in mind the construction of  
 145 knowledge bases for event recognition applications. These applications [5] process sequences  
 146 of simple events, such as sensor data, and recognize complex events of interest, i.e. events  
 147 that satisfy some pattern. Logic-based event recognition typically uses a knowledge base of  
 148 first-order rules to represent complex event patterns and a reasoning engine to detect such  
 149 patterns in the incoming data. In OLED this knowledge base is in the form of domain-specific  
 150 axioms in the Event Calculus, i.e. rules that specify the conditions under which simple,  
 151 low-level events initiate or terminate complex events.

152 OLED is using an online (single-pass) learning strategy. Online machine learning is a  
 153 method of machine learning in which data becomes available in a sequential order and is  
 154 used to update our best predictor for future data at each step, as opposed to batch learning  
 155 techniques which generate the best predictor by learning on the entire training data set  
 156 at once. To manage it, the Hoeffding bound [9] for evaluating clauses on a subset of the  
 157 input stream, is used. With this approach, significant speed-ups are obtained in training  
 time. Table 4 presents an example of input data that is provided to OLED. It consists of a

■ **Table 3** Theory Learnt by OLED

$\text{initiatedAt}(\text{growth}(X0),T) \leftarrow$ $\text{happensAt}(\text{size}(X0,3),T),$ $\text{happensAt}(\text{density}(X0,4),T).$
---

158 narrative and an annotation list. Narratives are the simple (or low level) events in terms of  
 159  $\text{happensAt}/2$ , expressing the values of communities' features. i.e.  $\text{happensAt}(\text{size}(c1,3),1)$ .  
 160 denotes that community  $c1$  has size 3 in time 1. Annotations are the complex (or high level  
 161 events) events in terms of  $\text{holdsAt}/2$ , expressing the ground truth for our training set. i.e.  
 162  $\text{holdsAt}(\text{growth}(c1),2)$ . denotes that community  $c1$  grew in time 2. The non-existence of  
 163  $c1$ 's annotation in time 1 states that growth event is terminated in time 1. Table 3 shows  
 164 an example of the theory OLED learnt after training. It represents we will begin to have  
 165 a growth event in time  $T+1$  for any community, which has size 3 and density 4 in time  $T$ .  
 166

167 This rule extracted because with these community features in time 1, we had a growth event  
 168 in time 2 (Table 4).

■ **Table 4** Input of OLED

Timeframe 1	Timeframe 2
<u>Narrative</u> happensAt(size(c1,3),1). happensAt(density(c1,4),1).	<u>Narrative</u> happensAt(size(c1,5),2). happensAt(density(c1,5),2).
<u>Annotation</u>	<u>Annotation</u> holdsAt( <i>growth</i> (c1),2).

168

169 **Learning** OLED learns a clause in a top-down fashion, by gradually adding literals to its  
 170 body. Instead of evaluating each candidate specialization on the entire input, it accumulates  
 171 training data from the stream, until the Hoeffding bound allows to select the best specializa-  
 172 tion. The instances used to make this decision are not stored or reprocessed but discarded  
 173 as soon as OLED extracts from them the necessary statistics for clause evaluation.

174 OLED relaxes the Lfi requirement that a hypothesis  $H$  covers every training interpretation,  
 175 and thus seeks a theory with a good fit in the training data. Let  $B$  consist of the domain-  
 176 independent EC axioms,  $r$  be a clause and  $I$  an interpretation. We denote by  $narrative(I)$   
 177 and  $annotation(I)$  the narrative and the annotation part of  $I$  respectively (Table 4). We  
 178 denote by  $Mr_I^r$  an answer set of  $B \cup narrative(I) \cup r$ . Given an annotation atom  $\alpha$  we say  
 179 that:

- 180 ■  $\alpha$  is a true positive (TP) atom clause  $r$ , iff  $\alpha \in annotation(I) \cap Mr_I^r$ .
- 181 ■  $\alpha$  is a false positive (FP) atom clause  $r$ , iff  $\alpha \in Mr_I^r$  but  $\alpha \notin annotation(I)$ .
- 182 ■  $\alpha$  is a false negative (FN) atom clause  $r$ , iff  $\alpha \in annotation(I)$  but  $\alpha \notin Mr_I^r$ .

183 We define a heuristic clause evaluation function  $G$  as follows:

$$184 \quad G(r) = \begin{cases} \frac{TP_r}{TP_r + FP_r}, & \text{if } r \text{ is an initiatedAt clause} \\ \frac{TP_r}{TP_r + FN_r}, & \text{if } r \text{ is a terminatedAt clause} \end{cases}$$

185 where  $TP_r$ ,  $FP_r$  and  $FN_r$  are the accumulated TP, FP and FN counts of clause  $r$  over the  
 186 input stream and  $G \in [0, 1]$ .

187 On the arrival of new interpretations, OLED either expands  $H$ , by generating a new  
 188 clause, or tries to expand (specialize) an existing clause. Clauses of low quality are pruned,  
 189 after they have been evaluated on a sufficient number of examples. Below there is an example  
 190 of OLED execution. Initially, processes Linit and Lterm start with two empty hypotheses,  
 191 Hinit and Hterm. Assume that the annotation in one of the incoming interpretations dictates  
 192 that the *growth* complex event holds at time 10, while it does not hold at time 9. Since  
 193 no clause in Hinit yet exists to initiate *growth* at time 9, Linit detects the *growth* instance  
 194 at time 10 as a FN and proceeds to theory expansion, generating an initiation clause for  
 195 *growth*. Lterm is not concerned with initiation conditions, so it will take no actions in this  
 196 case. Then, a new interpretation arrives, where the annotation dictates that *growth* holds  
 197 at time 20 but does not hold at time 21. In this case, since no clause yet exists in Hterm  
 198 to terminate *growth* at time 20, Lterm will detect an FP instance at time 21. It will then  
 199 proceed to theory expansion, generating a new termination condition for *growth*. At the same  
 200 time, assume that the initiation clause in Hinit is over-general and erroneously re-initiates  
 201 *growth* at time 20, generating an FP instance for the Linit process at time 21. In response to

202 that, Linit will proceed to clause expansion, penalizing the over-general initiation clause by  
 203 increasing its FP count, thus contributing towards its potential replacement by one of its  
 204 specializations.

205 In EC the initiation/termination of complex events depends only on the simple events  
 206 and contextual information of the previous time-point, therefore each interpretation is an  
 207 independent training instance. This guarantees the independence of observations that is  
 208 necessary for using the Hoeffding bound, which is defined in Appendix [A]. Let  $r$  be a clause  
 209 and  $G$  a clause evaluation function. Assume also that after  $n$  training instances,  $r1$  is  $r$ 's  
 210 specialization with the highest observed mean G-score  $\overline{G}$  and  $r2$  is the second best one, i.e.  
 211  $\Delta\overline{G} = \overline{G}(r1) - \overline{G}(r2) > 0$ . Then by the Hoeffding bound we have that for the true mean of  
 212 the scores' difference  $\Delta\hat{G}$  it holds  $\Delta\hat{G} > \Delta\overline{G} - \varepsilon$ , with probability  $1 - \delta$ . Hence, if  $\Delta\overline{G} > \varepsilon$   
 213 then  $\Delta\hat{G} > 0$ , implying that  $r1$  is indeed the best specialization to select at this point,  
 214 with probability  $1 - \delta$ . In order to decide which specialization to select, it thus suffices to  
 215 accumulate observations from the input stream until  $\Delta\overline{G} > \varepsilon$ . Also, because OLEP allows to  
 216 build decision models using only a small subset of the data, by relating the size of this subset  
 217 to a user-defined confidence level on the error margin of not making a (globally) optimal  
 218 decision, manages to consume small amounts of memory and time resources.

## 219 4 Proposed Methodology

220 A dynamic social network is time-tamped, and to be analysed it is *segmented* into time  
 221 frames, with an overlap between them to allow for a smooth transition. The problem we are  
 222 addressing is to predict the form of a community in the next frame, given some features of  
 223 the existing form of a community. The model that perform the prediction is learnt through  
 224 ILP and represented as clauses of Event Calculus.

### 225 4.1 Community Features

226 Pavlopoulou et al. [15] designed two types of features, the *structural* and the *temporal* ones.  
 227 Structural features represented the physical characteristics of a community such as size,  
 228 density etc. The temporal features included structural features and evolutionary events  
 229 that were derived from the past instances of a community, and from relations between past  
 230 instances of a community. In this work, we use the same features which describe below but  
 231 first let us introduce some notation:  $C_t$  is the set of communities at time frame  $F_t$ ;  $C_t^k$  is the  
 232 community  $k$  of set  $C_t$ ;  $n(F_t) = |V_t|$  is the size of set  $V_t$ ;  $m(F_t) = |E_t|$  is the size of set  $E_t$ ;  
 233 and  $C_{t_j}^{k_j}$  is the ancestor of  $C_{t_i}^{k_i}$ , where  $j < i$ .

234 The ancestors of a community do not necessarily belong to consecutive time frames. In  
 235 the current experiments, each community is tracked in each timeframe until its dissolution,  
 236 thus there are situations in which a community disappears at timeframe  $t_i$  but it reappears at  
 237 timeframe  $t_j$ , where  $j - i > 1$ . Thus, the  $i$ -th ancestor of a community is the  $i$ -th appearance  
 238 of the community in the past, counting from the present. Next are the features we used in  
 239 detail:

#### 240 Structural Features

241 **Size** is the normalized value for the size of a community  $C_t^k$  in time frame  $F_t$ :

$$242 \text{Size}(C_t^k) = \frac{n(C_t^k)}{n(F_t)}$$

243 **Density** is the number of  $C_t^k$  edges to the maximum number of edges the community  
244 could have:

$$245 \quad \text{Density}(C_t^k) = \frac{m(C_t^k)}{n(C_t^k)(n(C_t^k) - 1)/2}$$

246 **Cohesion** is defined as:

$$247 \quad \text{Cohesion}(C_t^k) = \frac{2m(C_t^k)(n(F_t) - n(C_t^k))}{m_{out}(C_t^k)(n(C_t^k) - 1)}$$

248 **Normalised Association** is defined as:

$$249 \quad \text{NormalizedAssociation}(C_t^k) = \frac{2m(C_t^k)}{2m(C_t^k) + m_{out}(C_t^k)}$$

250 **Ratio Association** is the average internal degree of a community's members:

$$251 \quad \text{RatioAssociation}(C_t^k) = \frac{2m(C_t^k)}{n(C_t^k)}$$

252 **Ratio Cut** is the average external degree of a community's members:

$$253 \quad \text{RatioCut}(C_t^k) = \frac{m_{out}(C_t^k)}{n(C_t^k)}$$

254 **Normalized Edges Number** is defined as:

$$255 \quad \text{NormalizedEdgesNumber}(C_t^k) = \frac{m(C_t^k)}{m(F_t)}$$

256 **Average Path Length** shows how close on average two random nodes are:

$$257 \quad \text{AveragePathLength}(C_t^k) = \frac{\sum_{v,u \in V_t^k, v \neq u} \text{dist}(v, u)}{n(C_t^k)(n(C_t^k) - 1)}$$

258 where  $\text{dist}(v, u)$  indicates the shortest distance between nodes  $v$  and  $u$ .

259 **Diameter** is the maximum shortest path between all pairs of nodes in community  $C_t^k$ :

$$260 \quad \text{Diameter}(C_t^k) = \max_{u,v \in V_t^k, u \neq v} \text{dist}(u, v)$$

261 **Clustering Coefficient** of community  $C_t^k$  shows how often, on average, the neighbours  
262 of a node of the community are also connected to each other.

263 **Centrality** measures how central each node of a community  $C_t^k$  is. We used three central-  
264 ity measures as features, namely closeness, betweenness and eigenvector centrality [17].

## 265 Temporal features

### 266 Structural features and Evolutionary events of N ancestors:

267 One group of temporal features is all the structural features, as described above, as well  
268 as the *evolutionary events* for the first  $n$  immediate ancestors of community  $C_{t_p}^{k_p}$ .

269 Another group of temporal features concerns pairs of communities and depict how a  
270 community has evolved compared to its previous instance in time. Using these pairs of  
271 communities for a given number of ancestors  $n$  to use, we compute the following temporal  
272 features:

273 **Similarity of consecutive communities** is the fraction between of nodes/edges that  
 274 are common in both instances of the community and total nodes/edges of two instances.

$$275 \quad JaccNodes(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|V_{t_i}^{k_i} \cap V_{t_{i-1}}^{k_{i-1}}|}{|V_{t_i}^{k_i} \cup V_{t_{i-1}}^{k_{i-1}}|}, JaccEdges(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|E_{t_i}^{k_i} \cap E_{t_{i-1}}^{k_{i-1}}|}{|E_{t_i}^{k_i} \cup E_{t_{i-1}}^{k_{i-1}}|},$$

$$276 \quad JaccNodes\&Edges(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|V_{t_i}^{k_i} \cap V_{t_{i-1}}^{k_{i-1}}| + |E_{t_i}^{k_i} \cap E_{t_{i-1}}^{k_{i-1}}|}{|V_{t_i}^{k_i} \cup V_{t_{i-1}}^{k_{i-1}}| + |E_{t_i}^{k_i} \cup E_{t_{i-1}}^{k_{i-1}}|}$$

278 where  $E_{t_i}^{k_i}$  is the set of edges and  $V_{t_i}^{k_i}$  the set of nodes of community  $C_{t_i}^{k_i}$ .

279 **Join nodes ratio** is the percentage of nodes joining the dynamic community as it evolves.

$$280 \quad JoinNodesRatio(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|V_{t_i}^{k_i} \setminus V_{t_{i-1}}^{k_{i-1}}|}{|V_{t_i}^{k_i}|}$$

281 **Left nodes ratio** is the percentage of nodes leaving the dynamic community as it evolves.

$$282 \quad LeftNodesRatio(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|V_{t_{i-1}}^{k_{i-1}} \setminus V_{t_i}^{k_i}|}{|V_{t_{i-1}}^{k_{i-1}}|}$$

283 **Activeness** is the ratio of the number of edges in current community  $C_{t_i}^{k_i}$  which also  
 284 existed in its ancestor community  $C_{t_{i-1}}^{k_{i-1}}$ , to the number of nodes in current community  
 285  $C_{t_i}^{k_i}$ .

$$286 \quad Activeness(C_{t_i}^{k_i}, C_{t_{i-1}}^{k_{i-1}}) = \frac{|E_{t_i}^{k_i}| - |E_{t_i}^{k_i} \setminus E_{t_{i-1}}^{k_{i-1}}|}{|V_{t_i}^{k_i}|}$$

287 The last two temporal features are computed for individual communities instead of pairs.

288 **Lifespan** is the ratio of the ancestors of a community based on the corresponding dynamic  
 289 community  $M$ , to the maximum number of ancestors it could have.

$$290 \quad LifeSpan(C_{t_w}^{k_w}) = \frac{|\{C_{t_p}^{k_p} \in M : p < w\}|}{t_w - 1}$$

291 **Aging** of a community  $C_{t_w}^{k_w}$  is the average age of the community members. The age of  
 292 a member is increased by 1 every time it is found to be also a member of an ancestor  
 293 community of  $C_{t_w}^{k_w}$  in the corresponding dynamic community. Aging is normalized by  
 294 dividing with the maximum possible age of members, which equals  $w$ .

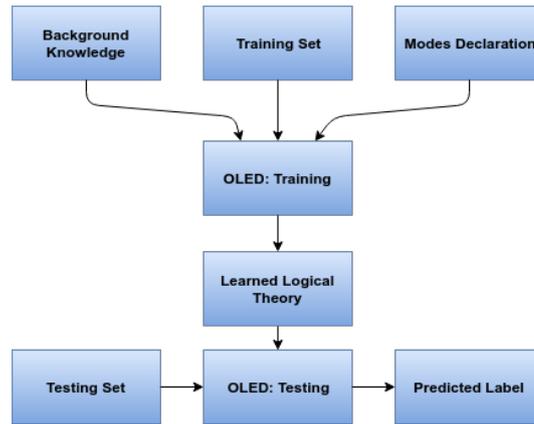
$$295 \quad Aging(C_{t_w}^{k_w}) = \frac{\sum_{v \in V_{t_w}^{k_w}} |\{C_{t_p}^{k_p} \in M : p \leq w, v \in V_{t_p}^{k_p}\}|}{(|\{C_{t_p}^{k_p} \in M : p < w\}| + 1)n(C_{t_w}^{k_w})}$$

296 **Feature Quantization** In OLED the values of variables are discrete thus we implemented  
 297 two methods to quantize variables. Let  $q_{value}$  be the number of quantized values,  $f_v$  be  
 298 the set with values of feature  $f$ . In first method, for each feature we split values' total  
 299 range to  $q_{value}$  intervals and the width of each is  $\frac{\max\{f_v\} - \min\{f_v\}}{q_{value}}$ . Thus the first interval is  
 300  $(\min\{f_v\}, \min\{f_v\} + q_{value})$ , the second is  $(\min\{f_v\} + q_{value}, \min\{f_v\} + 2q_{value})$  and so  
 301 on. The quantized value of each feature is the index of the interval it belongs to. The second  
 302 method sorts feature's values in a list and creates  $q_{value}$  sets. Taking one by one the values  
 303 from sorted list, we begin to fill the  $q_{value}$  sets with consecutive values until each set has  
 304  $\frac{|f_v|}{q_{value}}$  feature's values. Finally, if at least one feature or tag of community  $C_k$  is missing we  
 305 delete the  $C_k$ .

306 **4.2 Community Evolution Prediction**

307 OLED was used to predict four evolutionary events: *growth, shrinkage, continuation, dissolu-*  
 308 *tion*. Note that OLED handles two-class problems, so it predicts if a community will sustain  
 309 or will stop sustaining an evolutionary event. In Figure 1 we present the architecture of the  
 prediction system.

■ **Figure 1** Learning Architecture



310 The performance of an ILP system may degrade if the background knowledge provided  
 311 contains large amounts of irrelevant information so experts are required to set the background  
 312 knowledge they believe to be useful. Table 5 presents an example of background knowledge,  
 313 where the following types of rules can be defined: Rules for community entity recognition;  
 314 rules for time entity recognition; facts for features' quantized values recognition; rules for  
 315 values of ground truth recognition; and rules which represent the inertia of Event Calculus.  
 316 OLED can produce predicates of many forms. For example, an argument of a predicate can  
 317 be considered as input or as output. *Modes* declaration is a language that limits the forms a  
 318 predicate can have. Table 7 presents an example of modes.  
 319

■ **Table 5** Example of A OLED's Background Knowledge File

Background Knowledge File	
holdsAt( $F, T_e$ ) :- fluent( $F$ ), holdsAt( $F, T_s$ ), not terminatedAt( $F, T_s$ ), $T_e = T_s + 1$ , time( $T_s$ ),time( $T_e$ ).	holdsAt( $F, T_e$ ) :- fluent( $F$ ), initiatedAt( $F, T_s$ ), $T_e = T_s + 1$ , time( $T_s$ ),time( $T_e$ ). Inertia of Event Calculus
fluent( <i>growth</i> ( $X$ )) :- community( $X$ ).	Ground truth recognition
community( $X$ ) :- happensAt(size( $X, \_$ ), $\_$ ). community( $X$ ) :- happensAt(density( $X, \_$ ), $\_$ ).	Community entity recognition
time( $X$ ) :- happensAt(size( $\_$ , $\_$ ), $X$ ). time( $X$ ) :- happensAt(density( $\_$ , $\_$ ), $X$ ).	Time entity recognition
value(1..5).	Features' quantized values recognition

320 The form of rules that OLED learns is presented in Table 6. In the head of the rule,  
 321 *predicted\_event* is one of labels we try to predict (*growth, shrinkage, continuation, dissolu-*

■ **Table 6** Rules That OLED Learns

Rules
$\begin{aligned} &\text{initiatedAt/terminatedAt}(\langle \text{predicted\_event} \rangle(\langle \text{community}_i \rangle), \langle \text{time}_j \rangle) :- \\ &\quad \text{happensAt}(\langle \text{feature}_1 \rangle(\langle \text{community}_i \rangle, \langle \text{value}_1 \rangle), \langle \text{time}_j \rangle), \\ &\quad \dots \\ &\quad \text{happensAt}(\langle \text{feature}_n \rangle(\langle \text{community}_i \rangle, \langle \text{value}_n \rangle), \langle \text{time}_j \rangle). \quad (1)/(2) \end{aligned}$

■ **Table 7** Example of A OLED’s Mode Declarations File

Mode Declarations File	
<code>modeh(initiatedAt(<i>growth</i>(+community),+time))</code>	
<code>modeh(terminatedAt(<i>growth</i>(+community),+time))</code>	The form of the rule’s head
<code>modeb(happensAt(size(+community,#value),+time))</code>	
<code>modeb(happensAt(density(+community,#value),+time))</code>	The form of the rule’s body

322 *tion*). Notice that the  $community_i$  and  $time_j$  indices are the same in the body and head.  
 323 Rules can be interpreted as if  $feature_1$  of community  $community_i$  has  $value_1$  at  $time_j$   
 324 and the same is true for the rest of features then the initiation of event  $predicted\_event$  is  
 325 fired. This means that the  $predicted\_event$  will start to occur at  $time_{j+1}$ . The  $happensAt$   
 326 predicates that are required will be discovered by OLED. Likewise, if the body of rule (2) is  
 327 true then the termination of event  $predicted\_event$  is fired, thus the  $predicted\_event$  will  
 328 stop to occur at time  $time_{j+1}$ .

329 **Training and Testing** As shown in Figure 1, the dataset was split into training and testing  
 330 sets according to the Time Series Cross Validation, because it takes into account the temporal  
 331 relationship between the training and testing sets. Each training comprises only observations  
 332 that occurred prior to the observation.

**Fold 1:** Training set includes low events of communities from timeframes  $F_1, F_2$  and high events of communities in timeframe  $F_2$ . Testing set includes low events of communities from timeframe  $F_2$  and high events of communities from timeframes  $F_2, F_3$ .

**Fold 2:** Training set includes low events of communities from timeframes  $F_1, F_2, F_3$  and high events of communities from timeframe  $F_2, F_3$ . Testing set includes low events of communities from timeframe  $F_3$  and high events of communities from timeframes  $F_3, F_4$ .

...

**Fold T-2:** Training set includes low events of communities from timeframes  $F_1, F_2, \dots, F_{T-1}$  and high events of communities from timeframe  $F_2, F_3, \dots, F_{T-1}$ . Testing set includes low events of communities from timeframe  $F_{T-1}$  and high events of communities from timeframes  $F_{T-1}, F_T$ .

334  $T$  is total number of timeframes in the dataset. Note that the first timeframe has no  
 335 evolutionary events (high events) since there is no previous timeframe in order to track  
 336 the communities’ evolution of the first timeframe. Respectively, the last timeframe has  
 337 not features (low events) because there is no next timeframe to predict evolution of its  
 338 communities. Also, in the training set we comprise the low level events of the timeframe  
 339 that we are going to predict so that OLED extracts the time variable for high level events.  
 340 Finally, notice that in the testing set we also comprise the high level events of the previous  
 341 timeframe than that we are going to predict. This is required by OLED to initiate the inertia

342 of every community's event.

343 OLED as an online learner splits its input into chunks. In our experiments, we choose  
344 chunks of size 2. Thus, the imported timeframes for the training procedure are split into  
345 chunks two by two. We changed the functionality of OLED so that it creates rolling chunks.  
346 It means that first chunk contains the timeframes 1,2; the second one the timeframes 2,3;  
347 the third the timeframes 3,4 and so on. This necessary because, for example, timeframe 2  
348 has to be in the first and second chunk. In the first chunk, we need the high level events of  
349 timeframe 2 for getting the ground truth. While in the second chunk we use the low level  
350 events of timeframe 2 as features for our supervised learning classification.

351 The outline of training process is the following: Initially there is an empty theory. Each  
352 time OLED receives a chunk of training examples and transforms the existing theory to  
353 satisfy as close as possible the right prediction of the current examples. When the training  
354 process is completed, a logical theory is derived as the learnt model. Its form is illustrated in  
355 Table 6. Using this theory we predict the evolutionary events of communities which are in  
356 testing set.

## 357 5 Experiments

358 **Dataset description** The data were collected from the Mathematics Stack Exchange forum<sup>1</sup>,  
359 which is a question and answer site for mathematics. All questions are tagged with their  
360 subject areas. The dataset comprises 376,030 posts, 261,600 answers and comments, between  
361 28-09-2009 and 31-05-2013. Each user is represented by a node in a graph and there is an  
362 edge between two user nodes if one of them posts an answer or a comment on the other  
363 user's post. The dataset was split into 10 equally sized, with respect to the number of posts  
364 (questions, answers or comments), timeframes with 60% overlap between them.

365 Building the ground truth means obtaining community labels per time frame, and then  
366 obtaining the evolution of each community across time frames. We considered that a group  
367 of users belongs in the same community if they post (questions, answers or comments) about  
368 the same topic. In particular, we used tags to determine the communities and since on each  
369 post there are multiple tags, thus each user will be assigned to multiple communities. Answer  
370 and comment posts inherit the tag of the question they correspond to. Also communities  
371 with no more than 3 members were removed. The evolutionary events of each community  
372 (*Growth*, *Shrinkage*, *Continuation*, *Dissolution*) were obtained by thresholding. In particular  
373 if the size of the community in the next time frame is more (less) than 30 nodes compared to  
374 the size in the current frame then the community grows (shrinks).

375 There are communities which do not appear in each timeframe, although they may not  
376 have been dissolved yet. It happens because communities with few members in a timeframe  
377 are pruned from dataset. So, we are looking for the evolution of a community in every  
378 timeframe of the dataset and consider a community as dissolved only after its last appearance.  
379 The evolutionary events of dataset are imbalanced. In particular, the percentage of each  
380 class is: *Growth*: 0.5%, *Shrinkage*: 0.2%, *Continuation*: 90% and *Dissolution*: 0.3%. The  
381 features were quantized into 5 levels, and represented as low level events. The high level  
382 events are the evolutionary events. Experiments were executed with both structural and  
383 temporal features, where the number of ancestors was set to 4. At the end, the dataset  
384 was split in training and testing sets using Time Series Cross Validation method. Because

---

<sup>1</sup> <https://archive.org/download/stackexchange>

## 23:12 Predicting the Evolution of Communities with Online Inductive Logic Programming

385 the data are highly imbalanced apart from Micro Average measures, we also used Macro  
386 Averages.

387 **Survival Experiment** First, we conducted an experiment with an event, named *survival*  
388 that incorporated all the *growth*, *shrinkage* and *continuation* events. We used both structural  
and temporal features. In Table 8 we present the results. The micro measures are very high

■ **Table 8** Survival Experiment

	Survival Structural	Survival Temporal
Micro/Macro Precision	0.9737/0.8125	0.9882/0.9941
Micro/Macro Recall	0.9949/0.6289	1.0000/0.6765
Micro/Macro Fscore	0.9842/0.7090	0.9941/0.8051

389 because the survival events are 97% of the total data. As OLE initialized the inertia of the  
390 survival event, it did not find enough negatives examples (dissolution events) to fail in its  
391 prediction. Thus the Macro values are much lower.  
392

■ **Table 9** All events Structural features

	Growth	Shrinkage	Continuation	Dissolution
Micro/Macro Precision	0.2358/0.6037	0.1884/0.5832	0.9293/0.8066	0.6512/0.8125
Micro/Macro Recall	0.3027/0.6316	0.1512/0.5671	0.9760/0.6939	0.2629/0.6289
Micro/Macro Fscore	0.2651/0.6174	0.1677/0.5750	0.9521/0.7460	0.3746/0.7090

393 **Experiment with all events** The results on all events with *structural* features appear in  
Table 9. The macro precision is highest in the *dissolution*. The dataset with the *temporal*

■ **Table 10** All events Temporal features

	Growth	Shrinkage	Continuation
Micro/Macro Precision	0.1828/0.5730	0.1882/0.5743	0.9182/0.9222
Micro/Macro Recall	0.1828/0.5730	0.1633/0.5649	0.9955/0.6932
Micro/Macro Fscore	0.1828/0.5730	0.1749/0.5696	0.9553/0.7915

394 *features* contains the features of the previous 4 instances of a community, the first timeframe  
395 for this dataset is at time 5 (see Table 10). The theory which was derived for the dissolution  
396 event was empty. The predictor could not evaluate any rule with high score because there  
397 were not many available examples, since the number of timeframes (6, from  $F_5$  to  $F_{10}$ )  
398 and the number of communities is small. Thus, the dissolution event is not included in the  
399 experiments with *temporal features*. The *growth* and *shrinkage* events with temporal features  
400 have lower performance than the best corresponding events with *structural* features. But for  
401 the *continuation* event, the reverse is true for both micro and macro values.  
402

403 **Experiment with pruning** A learnt theory can be pruned to remove clauses whose score  
404 is smaller than a quality threshold  $S_{min}$ . In the previous experiments  $S_{min}$  was 0.9. Next  
405 we tried for each event the values 0.5, 0.7, 0.3 as  $S_{min}$  and choose the ones with the best  
406 performance. With the structural features, the best pruning value for the *growth* event is

■ **Table 11** Best Pruning Experiment with Structural Features

	Growth	Shrinkage	Continuation	Dissolution
Micro/Macro Precision	0.2343/0.6035	0.2047/0.5913	0.9293/0.8066	0.6512/0.8125
Micro/Macro Recall	0.3295/0.6431	0.1512/0.5679	0.9760/0.6939	0.2629/0.6289
Micro/Macro Fscore	0.2739/0.6227	0.1739/0.5794	0.9521/0.7460	0.3746/0.7090

407 0.7, for *shrinkage* 0.5, for *continuation* 0.9 and for *dissolution* 0.9 (see Table 11). In the  
 408 temporal features, the best pruning value for the *growth* event is 0.7, for *shrinkage* 0.7, and  
 409 for *continuation* 0.9 (see Table 12). Overall, pruning had a marginal improvement on the  
 results.

■ **Table 12** Best Pruning Experiment with Temporal Features

	Growth	Shrinkage	Continuation
Micro/Macro Precision	0.1828/0.5730	0.1951/0.5778	0.9182/0.9222
Micro/Macro Recall	0.1828/0.5730	0.1633/0.5656	0.9955/0.6932
Micro/Macro Fscore	0.1828/0.5730	0.1778/0.5716	0.9553/0.7915

410

411 **Experiment with long range rules** We changed the way rules are formed so that they  
 412 contain features of any of a community’s ancestors. This is similar to the temporal features,  
 413 but we do not have their values as different features but as the same features at different time  
 414 steps. In order to change the form of the derived rules we changed the modes declaration  
 415 so that the new form of rules captures long range relationships (see also Section 4.2). The  
 416 form of the new rules is shown in Table 13, which is the same as in Table 6 except of the  
 417  $\langle time \rangle$  value in the head can be different from that in the body and the  $geqn/3$  predicate,  
 418 which denotes that  $time_k$  is  $num_1$  units after  $time_l$ . Also because now OLED could include  
 419 more than two timeframes we had to increase the chunk size. If the chunk size equals to  
 420  $N + 2$ , then the derived rules can contain up to  $N$  ancestors’ features. However, a big chunk  
 421 size entails greater CPU and memory requirements. In the experiments we selected a chunk  
 422 of size 3, so we obtained rules that contained features of the first ancestor. The results are  
 423 presented in Table 14.

■ **Table 13** New Rules That OLED Learns With Long Range Relationships

Rules
<pre> initiatedAt/terminatedAt(<math>\langle predicted\_event \rangle</math>(<math>\langle community_i \rangle</math>),<math>\langle time_j \rangle</math>) :-   happensAt(<math>\langle feature_1 \rangle</math>(<math>\langle community_i \rangle</math>,<math>\langle value_1 \rangle</math>),<math>\langle time_1 \rangle</math>),   ...,   happensAt(<math>\langle feature_n \rangle</math>(<math>\langle community_i \rangle</math>,<math>\langle value_n \rangle</math>),<math>\langle time_n \rangle</math>),   geqn(<math>time_k</math>,<math>time_l</math>,<math>num_1</math>),   ...   geqn(<math>time_m</math>,<math>time_n</math>,<math>num_1</math>). </pre>

424 **Experiment with weighted TPs, FPs, FNs** Neither the previous method increased the  
 425 performance significantly. A problem is that the learnt theories contain more termination  
 426 than initiation rules, thus the initiation of some events does not happen. It means OLED

■ **Table 14** Long Range Relationships Experiment

	<b>Growth</b>	<b>Shrinkage</b>
<b>Micro/Macro Precision</b>	0.2446/0.6090	0.2016/0.5898
<b>Micro/Macro Recall</b>	0.3487/0.6527	0.1512/0.5678
<b>Micro/Macro Fscore</b>	0.2875/0.6300	0.1728/0.5786

427 predicts a negative event (i.e. event that does not occur) for a community at next timeframe  
 428 but in reality it is a positive event (i.e. the event occurs). In this case the FNs frequency  
 429 of OLED is increased. The numbers of initiation and termination rules are not balanced  
 430 because OLED evaluates its rules based on TPs, FPs and FNs values. Using these values, it  
 431 computes a score which evaluates the accuracy of a rule. To control score's value we can add  
 432 weights on TPs, FPs, FNs values during the training. For example, if the FNs weight is set  
 433 to 10, it means that the FNs will be considered as ten times more than it really is, in other  
 434 words the termination rules will overestimate the termination condition. Thus the score of  
 435 termination rules is getting decreased. With this way we focus more in quality than quantity  
 436 of termination rules. In Table 15, we present the best weights for each class in a experiment  
 with structural and temporal features. The results are presented in Table 16 and Table 17

■ **Table 15** Best weights for each class with structural/temporal features

	<b>TPs-weight</b>	<b>FPs-weight</b>	<b>FNs-weight</b>
<b>Growth</b>	1/1	1/5	15/1
<b>Shrinkage</b>	20/1	1/1	15/1
<b>Continuation</b>	1/1	1/1	1/15
<b>Dissolution</b>	1	1	15

437

for the experiment with structural features and with the temporal features respectively.

■ **Table 16** Weights on TPs,FPs,FNs - Experiment With Structural Features

	<b>Growth</b>	<b>Shrinkage</b>	<b>Continuation</b>	<b>Dissolution</b>
<b>Micro/Macro Precision</b>	0.2376/0.6055	0.1127/0.5533	0.9247/0.9623	0.8036/0.8878
<b>Micro/Macro Recall</b>	0.3487/0.6519	0.8023/0.8187	0.9845/0.6772	0.2113/0.6047
<b>Micro/Macro Fscore</b>	0.2826/0.6278	0.1977/0.6603	0.9537/0.7950	0.3346/0.7194

■ **Table 17** Weights on TPs,FPs,FNs - Experiment With Temporal Features

	<b>Growth</b>	<b>Shrinkage</b>	<b>Continuation</b>
<b>Micro/Macro Precision</b>	0.2184/0.5913	0.2459/0.6032	0.9182/0.9222
<b>Micro/Macro Recall</b>	0.2043/0.5857	0.1531/0.5654	0.9955/0.6933
<b>Micro/Macro Fscore</b>	0.2111/0.5885	0.1887/0.5837	0.9553/0.7915

438

While we were trying various values to weights, we noticed in the results that:

439

■ If TPs's weight is increased then TPs is increased, FPs is increased and FNs is decreased  
 441 because the number of initiations rules is increased.

442

■ If FPs's weight is increased then TPs is decreased, FPs is decreased and FNs is increased  
 443 because the number of initiations rules is decreased.

444 ■ If FNs's weight is increased then TPs is increased, FPs is increased and FNs is decreased  
445 because the number of termination rules is decreased.

446 We tried to increase the low TPs number by setting appropriate weights, but FPs also  
447 increased. OLED overestimated the initiation condition because its initiation rules are not  
448 specialised enough to detect correctly in which communities an event will occur. This is a  
449 strong indication that with the current features OLED performance could not improve. In  
450 Appendix A, we present some of the clauses that derived by above experiments.

## 451 **6** Conclusions

452 We tried to predict the evolution of communities in a dynamic social network. The evolution of  
453 a community is described as the occurrence of *growth*, *shrinkage*, *continuation* and *dissolution*  
454 events. We carried out the prediction using OLED, an Inductive Logic Programming  
455 system for learning logical theories from data streams. Initially, we tracked the evolution  
456 of communities over the time and obtained the ground truth of evolutionary events. As  
457 features we used structural characteristics of communities. Moreover, we also tried temporal  
458 features where a preset number of previous instances of each communities were used as well  
459 as features that capture change between consecutive instances of a community. Subsequently,  
460 the features were quantized. The dataset was obtained from the Mathematics Stack Exchange  
461 forum. We presented the micro and macro averages, because the classes (*Growth*, *Shrinkage*,  
462 *Continuation* and *Dissolution*) were unbalanced.

463 We also investigated the best pruning values for the theory in OLED, which did not  
464 improve the results. Overall, the experiments with the temporal features had a worse  
465 performance than the experiment with the structural features, probably because there were  
466 not many timeframes. Then, we execute experiments where OLED learnt rules that represent  
467 long range relationships between an evolutionary event and features.

468 Subsequently, weights were applied to TPs, FPs and FNs values to change rules' scores.  
469 This was the experiment with the best results. Finally, we presented the features that were  
470 the most influential for each evolutionary event.

471 Future work could be directed to a range of different fields. Others classifiers can be  
472 used to predict the evolution of communities (e.g. SVM, Random Forest) and compare them  
473 to our results. Additional, evolutionary events can be added such as merge or split. Other  
474 types of features (i.e. topics or context of discussions in social networks ) could be studied as  
475 well as features that capture the dynamics of communities, such as the rate of change of an  
476 existing feature. Also, another quantization algorithm, which will adapt better the quantized  
477 values to the distribution of the values of the features might help. Because OLED is an  
478 online system, it needs many data to decide if a rule is trusted. However, in our dataset we  
479 had only 10 frames. Thus, datasets with more timeframes could be examined. Moreover, a  
480 different segmentation of the data stream could be tried, to study its effect on the prediction.  
481 Finally, it would be very interesting to test the learnt rules in other datasets to notice how  
482 relevant they are at the problem of community evolution prediction.

## 483 **A** Appendix

484 An advantage of OLED is that the predictive model (Theory) it derives is human-readable.  
485 Thus the rules can be read, analyzed and interesting results can be derived from them. Some  
486 of the best performing rules are shown in Tables 18 and 19. Transferability to new datasets  
487 is also an interesting possibility.

■ **Table 18** Rules learnt in the best experiment: Growth and Shrinkage

<pre> initiatedAt(<i>growth</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(density(<i>X0</i>,1),<i>T1</i>),   happensAt(diameter(<i>X0</i>,2),<i>T1</i>). </pre>
<pre> terminatedAt(<i>growth</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(ratio_cut(<i>X0</i>,3),<i>T1</i>),   happensAt(average_path_length(<i>X0</i>,3),<i>T1</i>),   happensAt(normalized_edges_number(<i>X0</i>,5),<i>T1</i>). </pre>
<pre> terminatedAt(<i>growth</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(ratio_cut(<i>X0</i>,3),<i>T1</i>),   happensAt(closeness_centrality(<i>X0</i>,3),<i>T1</i>),   happensAt(normalized_edges_number(<i>X0</i>,5),<i>T1</i>). </pre>
<pre> terminatedAt(<i>growth</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(cohesion(<i>X0</i>,2),<i>T1</i>),   happensAt(average_path_length(<i>X0</i>,3),<i>T1</i>),   happensAt(diameter(<i>X0</i>,2),<i>T1</i>). </pre>
<pre> terminatedAt(<i>shrinkage</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(ratio_association(<i>X0</i>,3),<i>T1</i>). </pre>
<pre> terminatedAt(<i>shrinkage</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(average_path_length(<i>X0</i>,2),<i>T1</i>). </pre>
<pre> terminatedAt(<i>shrinkage</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(closeness_centrality(<i>X0</i>,2),<i>T1</i>),   happensAt(ratio_cut(<i>X0</i>,1),<i>T1</i>). </pre>
<pre> initiatedAt(<i>shrinkage</i>(<i>X0</i>),<i>T1</i>) :-   happensAt(eigenvector_centrality(<i>X0</i>,1),<i>T1</i>),   happensAt(ratio_association(<i>X0</i>,5),<i>T1</i>). </pre>

488 Some features appeared more often in rules of specific evolutionary events than others;  
489 while some never appeared. In Tables 20 and 21 we present for each evolutionary event  
490 (*growth*, *shrinkage*, *continuation*, *dissolution*), the frequency of the structural features the  
491 bodies of rules.

492 In Table 20 it can be noticed that features like *diameter*, *cohesion*, *ratio\_cut* and *aver-*  
493 *age\_path\_length* affect the prediction of the *growth* event since they represent 54.14% of total  
494 features which appeared in the rules. On the contrary features like *betweenness\_centrality*  
495 and *normalized\_association* did not appear at all. For the *shrinkage* event the most used  
496 features are *ratio\_association*, *ratio\_cut*, *cohesion* and *clustering\_coefficient*, while the  
497 *normalized\_association* feature does not appear. In Table 21 the features of the *continuation*  
498 and the *dissolution* events are presented. The continuation event seems to be affected  
499 mostly from *ratio\_cut*, *ratio\_association* and the *clustering\_coefficient* and not by the  
500 cohesion. While for the *dissolution* event, every feature is used in prediction, and especially  
501 the *clustering\_coefficient*, *cohesion* and the *betweenness\_centrality*.

502 In Tables 22 and 23 we present temporal features which appear in the rules of corresponding  
503 experiments. For the *growth* event the temporal features: *ancestor4\_average\_path\_length*, *ac-*  
504 *tiveness\_ancestor\_2\_ancestor3*, *aging\_ancestor0*, *ancestor1\_diameter*, *ancestor3\_closeness*  
505 *\_centrality* and the rest that are presented in Table 22, are the equally important. *Shrinkage*  
506 event prediction uses the values of *ancestor1\_event\_is\_shrinking*, *ancestor4\_clustering\_coe-*  
507 *fficient*, *cohesion*, *eigenvector\_centrality*, *joinNodesRatio\_currentCommunity\_ancestor0*,

■ **Table 19** Rules learnt: Survival

$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{ratio\_association}(X0,4),T1),$ $\text{happensAt}(\text{density}(X0,2),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{ratio\_association}(X0,3),T1),$ $\text{happensAt}(\text{clustering\_coefficient}(X0,4),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{diameter}(X0,3),T1),$ $\text{happensAt}(\text{ratio\_cut}(X0,3),T1),$ $\text{happensAt}(\text{ratio\_association}(X0,2),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{ratio\_association}(X0,3),T1),$ $\text{happensAt}(\text{closeness\_centrality}(X0,3),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{clustering\_coefficient}(X0,1),T1),$ $\text{happensAt}(\text{closeness\_centrality}(X0,5),T1),$ $\text{happensAt}(\text{cohesion}(X0,4),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{normalized\_edges\_number}(X0,2),T1),$ $\text{happensAt}(\text{cohesion}(X0,2),T1),$ $\text{happensAt}(\text{average\_path\_length}(X0,1),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{size}(X0,1),T1),$ $\text{happensAt}(\text{betweenness\_centrality}(X0,1),T1),$ $\text{happensAt}(\text{cohesion}(X0,3),T1).$
$\text{terminatedAt}(\text{survival}(X0),T1) :-$ $\text{happensAt}(\text{normalized\_edges\_number}(X0,1),T1),$ $\text{happensAt}(\text{cohesion}(X0,3),T1),$ $\text{happensAt}(\text{average\_path\_length}(X0,1),T1).$

508 *ratio\_cut*. Many temporal features are missing from the bodies of rules for both *growth* and  
 509 *shrinkage* events. For the *continuation* event only *cohesion* and *ratio\_cut* were used.

510 **Hoeffding bound** The *Hoeffding bound* is a statistical tool that is used as a probabilistic  
 511 estimator of the generalization error of a model (true expected error on the entire input),  
 512 given its empirical error (observed error on a training subset). Given a random variable  
 513  $X \in [0, 1]$  and an observed mean  $\bar{X}$  of its values after  $n$  independent observations, the  
 514 Hoeffding Bound states that, with probability  $1 - \delta$ , the true mean  $\hat{X}$  of the variable lies  
 515 in an interval  $(\bar{X} - \varepsilon, \bar{X} + \varepsilon)$ , where  $\varepsilon = \sqrt{\frac{\ln(1/\delta)}{2n}}$ . In other words, the true average can  
 516 be approximated by the observed one with probability  $1 - \delta$ , given an error margin  $\varepsilon$  that  
 517 decreases with the number of observations  $n$ .

■ **Table 20** Structural features frequency: Growth and Shrinkage

Growth	Percentage	Shrinkage	Percentage
diameter	17.68%	ratio_association	20%
cohesion	13.26%	ratio_cut	13.55%
ratio_cut	11.60%	cohesion	11.61%
average_path_length	11.60%	clustering_coefficient	10.97%
density	8.29%	eigenvector_centrality	9.03%
ratio_association	7.73%	density	8.39%
clustering_coefficient	7.73%	average_path_length	7.10%
size	6.63%	closeness_centrality	6.45%
closeness_centrality	6.08%	centrality	3.871%
eigenvector_centrality	3.87%	diameter	3.87%
normalized_edges_number	2.76%	betweenness_centrality	2.58%
centrality	2.76%	size	1.29%
		normalized_edges_number	1.29%

■ **Table 21** Structural features frequency: Continuation and Dissolution

Continuation	Percentage	Dissolution	Percentage
ratio_cut	16.07%	clustering_coefficient	25.93%
ratio_association	15%	cohesion	15.74%
clustering_coefficient	10%	betweenness_centrality	10.19%
density	9.64%	diameter	9.26%
diameter	8.21%	size	8.33%
closeness_centrality	8.21%	normalized_edges_number	6.48%
eigenvector_centrality	7.14%	ratio_association	5.56%
centrality	6.79%	closeness_centrality	3.70%
betweenness_centrality	6.43%	average_path_length	3.70%
normalized_association	4.64%	ratio_cut	2.78%
average_path_length	4.64%	normalized_association	2.78%
normalized_edges_number	2.5%	centrality	2.78%
size	0.71%	density	1.85%
		eigenvector_centrality	0.93%

■ **Table 22** Temporal features frequency: Growth

<b>Growth</b>	<b>Percentage</b>
ancestor4_average_path_length	12.5%
activeness_ancestor_2_ancestor3	6.25%
aging_ancestor0	6.25%
ancestor1_diameter	6.25%
ancestor3_closeness_centrality	6.25%
ancestor3_clustering_coefficient	6.25%
ancestor3_diameter	6.25%
ancestor4_centrality	6.25%
ancestor4_clustering_coefficient	6.25%
ancestor4_diameter	6.25%
jaccardCoefficient_ancestor_0_ancestor1	6.25%
jaccardCoefficient_ancestor_2_Ancessor3	6.25%
joinNodesRatio_ancestor_2_ancestor3	6.25%
joinNodesRatio_currentCommunity_ancestor0	6.25%
leftNodesRatio_ancestor_0_ancestor1	6.25%

■ **Table 23** Temporal features frequency: Shrinkage and Continuation

<b>Shrinkage</b>	<b>Percentage</b>
ancestor1_event_is_shrinking	16.66%
ancestor4_clustering_coefficient	16.66%
cohesion	16.66%
eigenvector_centrality	16.66%
joinNodesRatio_currentCommunity_ancestor0	16.66%
ratio_cut	16.66%
<b>Continuation</b>	<b>Percentage</b>
cohesion	50%
ratio_cut	50%

## 518 — References

- 519 **1** Hendrik Blockeel, Luc De Raedt, Nico Jacobs, and Bart Demoen. Scaling up induc-  
520 tive logic programming by learning from interpretations. *Data Mining and Knowledge*  
521 *Discovery*, 3(1):59–93, Mar 1999. URL: <https://doi.org/10.1023/A:1009867806624>,  
522 doi:10.1023/A:1009867806624.
- 523 **2** P. Bródka, P. Kazienko, and B. Kołoszczyk. Predicting Group Evolution in the Social  
524 Network. *ArXiv e-prints*, October 2012. arXiv:1210.5161.
- 525 **3** L. De Raedt. *Logical and Relational Learning*. Cognitive Technologies. Springer Berlin  
526 Heidelberg, 2008. URL: <https://books.google.gr/books?id=FFYIOXvwq7MC>.
- 527 **4** Georgios Diakidis, Despoina Karna, Dimitris Fasarakis-Hilliard, Dimitrios Vogiatzis, and  
528 George Paliouras. Predicting the evolution of communities in social networks. In *Pro-*  
529 *ceedings of the 5th International Conference on Web Intelligence, Mining and Semantics,*  
530 *WIMS '15*, pages 1:1–1:6, New York, NY, USA, 2015. ACM. URL: [http://doi.acm.org/](http://doi.acm.org/10.1145/2797115.2797119)  
531 [10.1145/2797115.2797119](http://doi.acm.org/10.1145/2797115.2797119), doi:10.1145/2797115.2797119.
- 532 **5** Opher Etzion and Peter Niblett. *Event Processing in Action*. Manning Publications Co.,  
533 Greenwich, CT, USA, 1st edition, 2010.
- 534 **6** Santo Fortunato. Community detection in graphs. *CoRR*, abs/0906.0612, 2009.
- 535 **7** Bogdan Gliwa, Piotr Bródka, Anna Zygumt, Stanislaw Saganowski, Przemyslaw Kazienko,  
536 and Jaroslaw Kozlak. Different approaches to community evolution prediction in blo-  
537 gosphere. In *Advances in Social Networks Analysis and Mining 2013, ASONAM '13,*  
538 *Niagara, ON, Canada - August 25 - 29, 2013*, pages 1291–1298, 2013. URL: [http:](http://doi.acm.org/10.1145/2492517.2500231)  
539 [//doi.acm.org/10.1145/2492517.2500231](http://doi.acm.org/10.1145/2492517.2500231), doi:10.1145/2492517.2500231.
- 540 **8** Mark Goldberg, Malik Magdon ismail, Srinivas Nambirajan, and James Thompson. Track-  
541 ing and predicting evolution of social communities. -, -.
- 542 **9** Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Jour-*  
543 *nal of the American Statistical Association*, 58(301):13–30, 1963. doi:10.1080/01621459.  
544 1963.10500830.
- 545 **10** Nagehan İlhan and Şule Gündüz Ögüdücü. Predicting community evolution based on  
546 time series modeling. In *Proceedings of the 2015 IEEE/ACM International Conference*  
547 *on Advances in Social Networks Analysis and Mining 2015, ASONAM '15*, pages 1509–  
548 1516, New York, NY, USA, 2015. ACM. URL: [http://doi.acm.org/10.1145/2808797.](http://doi.acm.org/10.1145/2808797.2808913)  
549 [2808913](http://doi.acm.org/10.1145/2808797.2808913), doi:10.1145/2808797.2808913.
- 550 **11** Sanjay Ram Kairam, Dan J. Wang, and Jure Leskovec. The life and death of online groups:  
551 Predicting group growth and longevity. In *Proceedings of the Fifth ACM International*  
552 *Conference on Web Search and Data Mining, WSDM '12*, pages 673–682, New York, NY,  
553 USA, 2012. ACM. URL: <http://doi.acm.org/10.1145/2124295.2124374>, doi:10.1145/  
554 2124295.2124374.
- 555 **12** Nikos Katzouris, Alexander Artikis, and Georgios Paliouras. Online learning of event  
556 definitions. *Theory and Practice of Logic Programming*, 16(5-6):817–833, 2016. doi:10.  
557 1017/S1471068416000260.
- 558 **13** Robert Kowalski and Marek Sergot. *A Logic-Based Calculus of Events*, pages 23–55.  
559 Springer Berlin Heidelberg, Berlin, Heidelberg, 1989. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-642-83397-7_2)  
560 [978-3-642-83397-7\\_2](https://doi.org/10.1007/978-3-642-83397-7_2), doi:10.1007/978-3-642-83397-7\_2.
- 561 **14** Akshay Patil, Juan Liu, and Jie Gao. Predicting group stability in online social networks.  
562 In *Proceedings of the 22Nd International Conference on World Wide Web, WWW '13*,  
563 pages 1021–1030, New York, NY, USA, 2013. ACM. URL: [http://doi.acm.org/10.1145/](http://doi.acm.org/10.1145/2488388.2488477)  
564 [2488388.2488477](http://doi.acm.org/10.1145/2488388.2488477), doi:10.1145/2488388.2488477.
- 565 **15** Maria Evangelia G. Pavlopoulou, Grigorios Tzortzis, Dimitrios Vogiatzis, and George  
566 Paliouras. Predicting the evolution of communities in social networks using structural  
567 and temporal features. In *12th International Workshop on Semantic and Social Media*

- 568 *Adaptation and Personalization, SMAP 2017, Bratislava, Slovakia, July 9-10, 2017*, pages  
569 40–45, 2017. URL: <https://doi.org/10.1109/SMAP.2017.8022665>, doi:10.1109/SMAP.  
570 2017.8022665.
- 571 **16** Mansoureh Takaffoli, Reihaneh Rabbany, and Osmar R. Zaiane. Community evolution  
572 prediction in dynamic social networks. doi:10.1109/ASONAM.2014.6921553.
- 573 **17** R. Zafarani, M.A. Abbasi, and H. Liu. *Social Media Mining: An Introduction*. Cambridge  
574 University Press, 2014. URL: <https://books.google.gr/books?id=H9FkAwwAAQBAJ>.