

MeetupNet Dublin: Discovering Communities in Dublin’s Meetup Network^{*}

Arjun Pakrashi, Elham Alghamdi, Brian Mac Namee, Derek Greene

University College Dublin, Ireland

arjun.pakrashi@ucdconnect.ie, elham.alghamdi@ucdconnect.ie,
derek.greene@ucd.ie, brian.macnamee@ucd.ie

Abstract. Meetup.com is a global online platform which facilitates the organisation of meetups in different parts of the world. A meetup group typically focuses on one specific topic of interest, such as sports, music, language, or technology. However, many users of this platform attend multiple meetups. On this basis, we can construct a co-membership network for a given location. This network encodes how pairs of meetups are connected to one another via common members. In this work we demonstrate that, by applying techniques from social network analysis to this type of representation, we can reveal the underlying meetup community structure, which is not immediately apparent from the platform’s website. Specifically, we map the landscape of Dublin’s meetup communities, to explore the interests and activities of meetup.com users in the city.

1 Introduction

Meetup.com is an online platform that helps people to plan, organise, and discover public or private events, referred to as *meetups*. Members of the platform can join specific meetup groups, that are usually based around a specific topic or activity, within which meetup events are organised. Meetup.com hosts groups that focus on diverse topics including sports, food, language, technology, business, philosophy, and dancing. The meetup.com platform is used worldwide and at the time of writing hosted more than 300k meetup groups and 39m users. On average over 3m people attend a meetup event each month.

A complex network structure underlies the meetup.com platform, consisting of connections between users, meetup groups, and meetup events. By applying popular techniques from the field of social network analysis, such as centrality analysis [5] or community finding [11], we can potentially reveal insights about a given city or country. This could range from the most active sports communities, to the most popular types of music, to the most vibrant tech scenes. These insights are potentially useful to meetup.com organisers and members, recruiters, entrepreneurs, tourists, and city planners.

In this paper, we describe an analysis of the meetup.com network in Dublin, Ireland. Rather than focusing on individual users of the platform, we instead

^{*} This research was supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289.

form a network representation at the meetup level, where a connection between two meetups exists if the two groups share members in common. We focus on two key research questions: 1) do distinct thematically-coherent communities exist within Dublin’s Meetup ecosphere?; 2) if so, how do these communities overlap with one another? To answer these questions, we apply the popular *OSLOM* (Order Statistics Local Optimisation Method) algorithm [15] to discover overlapping communities of connected meetups, and text analytics methods to summarise these communities. This analysis reveals Dublin’s place as a technology capital through a large number of communities of large technology focused meetups, the importance of meetups focusing on generic topics such as language learning as connectors within the community, and how large communities of meetups are linked by a small number of meetups. This analysis is a demonstration of how the meetup.com network structure can be used to gain insights into the communities that exist within a city and how they interact. To support the further analysis of meetup groups in other geographic locations, we make the relevant code and data available for reuse¹.

The remainder of this paper proceeds as follows. Section 2 describes related work on co-occurrence networks of the type used in this paper and community finding. Section 3 describes how the meetup network was constructed, how community finding techniques were applied to it, and how the resulting communities were labelled using textual metadata. Section 4 explores the network created and the communities found within it, focusing on a subset of largely technology-focused communities. Finally, Section 5 concludes with suggested directions for future work in this area.

2 Related Work

2.1 Analysing Co-occurrence Networks

A *co-occurrence network* is a weighted network constructed such that each edge indicates the frequency with which two items appear in the same context. In some cases, this kind of network is formed by projecting an existing bipartite network to a one-mode network by defining the weights as the number of co-occurrences. In other cases, such networks are created directly from raw count data. For instance, in a physical co-location network, we might create an edge between two individuals, where the edge weight indicates the number of times the individuals were in the same place at the same time. Network analysis techniques have been used to explore co-occurrence patterns in a range of domains. One common application area has been in bibliometrics, where co-citation networks have been analysed [12]. Here a co-citation link exists between two research papers if they are both cited by a third paper. By analysing the structure of such networks, it is possible to map the research activities, collaborations, and trends within and across research fields [20]. Other common examples of the use

¹ Code: <https://github.com/phoxis/MeetupNetDublin>
Interactive graph: <https://draig.ucd.ie/MeetupNetDublinInteractive/>

of co-occurrence networks include the analysis of word co-occurrence patterns in natural language processing [8], the analysis of co-purchasing trends in online retail systems [10], and the use of co-listed information to recommend users to follow in the context of online social media platforms [13].

2.2 Community Finding

When analysing networks in many domains, we will often be interested in performing community detection, where the goal is to identify the underlying group structures in the data. Typically this is performed as an unsupervised task. A substantial amount of work in this area has focused on the detection of disjoint communities, where each node belongs to at most one community. Algorithms in this context can be broadly grouped into three types:

(1) *Hierarchical algorithms* construct a tree of communities based on the network topology. These can be one of two types: divisive algorithms [11] or agglomerative algorithms [7]. (2) *Modularity-based algorithms* optimise the well-known modularity objective function to uncover communities in a network [17]. (3) *Other algorithms* which include those based on label propagation approaches [21], spectral methods that make use of the eigenvectors of a Laplacian or standard matrix, and methods based on statistical modelling [9].

In many real-world networks, we observe “pervasive overlap”, such that individuals frequently belong to many highly-overlapping communities [2]. Therefore, overlapping community finding algorithms [3] have been developed for application to these networks. These can be classified into four main categories:

(1) *Node seeding and local expansion algorithms* detect communities by starting from a node, or a small group of nodes, and then expanding these into a community using some fitness function. OSLOM [15] is one example, that expands communities based on a fitness function measuring the statistical significance of communities with respect to random variations. We discuss this algorithm in more detail in Section 3.3. (2) *Clique expansion methods* use a group of fully-connected nodes, called a clique, as the starting point for building larger communities. CFinder [1] and Greedy Clique Expansion (GCE) [16] are examples of this type of algorithm. (3) *Link clustering algorithms* detect communities by splitting the network edges rather than the nodes. GA-NET+ [18] is a good example of this category. (4) *Label propagation algorithms* attempt to group each node into a community based on its neighbouring nodes’ affinities. An example of this approach is the Speaker Label Propagation Algorithm (SLPA) [21].

3 Building the Dublin Meetup Network

This section describes how data was collected using the meetup.com API, how a meetup co-occurrence network was constructed, the application of community finding methods to this network, and the use of text analytics methods to label and explain communities found in the network.




Fig. 1: Example of a co-membership relationship between meetups, where an arrow indicates that a user is a member of a meetup.

3.1 Data Collection

The Meetup.com website provides an open API² that allows access to data from the meetup.com platform. Building a meetup co-occurrence network required data about meetup groups in Dublin and the users that are members of each group. The `/find/groups` API call generates a list of all meetups in a specified country. Metadata providing the name, description, and host city of each meetup is also returned. Using this call we generated a list of all meetups in Ireland and then filtered this to exclude those not hosted in Dublin. Private meetups were also excluded. This left a filtered set of 1,482 Dublin-based public meetups. The `/2/members` API call generates a list of member IDs for all members of a specific meetup group. We used this call to generate a list of all of the members of each Dublin-based public meetup group identified in the previous step.

3.2 Network Construction

Motivated by the concept of co-citation networks in bibliometrics, we create an undirected weighted graph based on the member overlaps, or *co-memberships*, between pairs of meetups. Fig. 1 illustrates this approach. As *Meetup 1* and *Meetup 2* share common users, a link would exist between them in a co-occurrence graph, whereas neither of these would link to *Meetup 3* as no common users exist.

In the proposed network, each unique meetup is represented as a node, and a weighted edge between two nodes represents an association between the two meetups represented by its endpoints. Every meetup has a set of registered members, which allows us to measure the degree of overlap between the member sets for pairs of meetups. Rather than looking at the raw number of common members, we use a normalised value. Specifically, the degree of association between two meetups is calculated using the Jaccard index [14], which has been previously used in co-citation analysis [19]. Formally, we calculate the edge weight w_{ij} between meetups i and j , with sets of members M_i and M_j respectively, as:

$$w_{ij} = \frac{|M_i \cap M_j|}{|M_i \cup M_j|} \quad (1)$$

² Meetup.com API: https://www.meetup.com/meetup_api/

Edges only exist between pairs of nodes for which $w_{ij} > 0$. We refer to the resulting network as the *normalised meetup co-membership network* or more simply the *Dublin meetup network*.

3.3 Finding Communities

We apply community finding to the Dublin meetup network to organise it into a smaller number of communities of related meetups that are easier to interpret, as opposed to manually inspecting a large number of meetups individually. Given that we might expect some users to be members of a diverse set of meetups, we apply an overlapping community finding approach to the co-membership network, which allows each meetup to potentially belong to multiple communities. Specifically, we apply the popular OSLOM algorithm [15], which is suitable for use on weighted networks.

OSLOM follows a greedy expansion strategy to detect communities by optimising a local fitness function on seed nodes based on statistical measurement. This is done in three stages. In the first stage, the algorithm searches for significant clusters by selecting a node at random as a seed then expanding it into a larger community. During this expansion, the algorithm tests the fitness score of the neighbouring nodes of the seed node by using a statistical test that evaluates the significance of each neighbouring node to be added to the seed. In the second stage, for each community, the algorithm performs a process which involves removing or adding nodes to maximise its significance score. In the third stage, the hierarchical levels of the communities are detected. Due to the presence of a stochastic element during calculating a node’s significant score the above stages are repeated several times and aggregated to create a stable set of communities.

To identify communities, we applied the undirected version of OSLOM to the meetup network using a range of values $[0.01, 0.5]$ for the resolution parameter, which indirectly controls the size of the communities identified by the algorithm. After each run, we filtered communities containing < 5 nodes, on the basis that these do not represent significant groupings of meetups. Based on manual inspection, a value of 0.1 for the resolution parameter provided a good trade-off between ensuring that communities were coherent, while also ensuring that the number of duplicate communities (*i.e.* related to identical topics) was minimised.

3.4 Labelling Communities

To extract meaningful insights from the community finding results, and to examine the topical coherence of these communities, we produce human-interpretable labels for each community. This allows us to explain and understand the groups at a high level. Fortunately, the meetup.com API provides a rich set of metadata which can be used for this purpose. Specifically, we propose a custom approach for labelling each community based on the short name field and the longer description field associated each meetup assigned to that community. Formally, we generate *name labels* for communities as follows:

1. For each meetup name field, extract all alphanumeric terms and filter out common stopwords (*e.g.* “the”, “meetup”, “group”).
2. Construct a meetup-term matrix A , such that each row corresponds to a meetup, each column corresponds to a term, and each entry indicates the number of times a term appears in a meetup name.
3. Apply standard log-based TF-IDF weighting to the matrix, and normalise the rows to unit length to account for different name lengths.
4. For each community C :
 - (a) From A , compute the mean vector of the rows corresponding to the meetups which have been assigned to C .
 - (b) Rank the values in the mean vector in descending order, and select the top t terms to create a name label.

We apply an analogous procedure to generate *description labels* for communities, based on the longer meetup description text.

4 Exploring the Dublin Meetup Network

In the original co-membership graph, each meetup group is represented by a node and the edges indicates the weight of the connection between pairs of meetups. The network consists of 1,482 nodes connected by 1,416,326 weighted edges, which represents a single connected component.

4.1 Network Characterisation

As is typical of many co-occurrence networks, the Dublin meetup network is highly dense, with edges present between 64.5% of all possible pairs of nodes. However, there is some variation in terms of the weights on these edges. We observe that 86.63% of the weighted edges have values ≤ 0.01 . This is indicative of a relatively small intersection between the memberships of many meetups.

When ranked by number of members, unsurprisingly, the three largest meetups are from the social sphere: *New and Not So New In Dublin* (21,149 members), *Events, Drinks and Dancing in Dublin* (16,582 members) and *Dublin International* (14,812 members). It is more informative, however, to measure the importance of meetup nodes in the overall meetup network and we do this using centrality analysis. A range of measures have previously been proposed for this task. We focus on two popular measures of centrality which are designed for use on weighted networks:

1. In *weighted eigenvector centrality*, a node is deemed more important if it is connected to other important nodes. In the weighted variant of this measure, centrality scores are calculated based on the first left eigenvector of the weighted graph adjacency matrix [5].
2. The *weighted betweenness centrality* measure identifies strategic bridges in a network. Nodes that occur on many shortest paths between other nodes in the network have high centrality. In the weighted variant of betweenness, the weighted distances between nodes are taken into account [6].

The top 10 meetups as ranked by each measure are listed in Table 1. It is interesting to note the significant difference between the lists generated by these two approaches—in fact there are no meetups that occur in both lists. The dominance of technology related meetups in Table 1a reflects the vibrance of the tech community in Ireland and indicates a slight bias on the meetup.com platform towards technology savvy users. It also suggests the existence of a cluster of similar meetups attended by a core group of overlapping members. The meetups in Table 1b reflect the fact that betweenness centrality measures the ability of nodes in a network to connect disparate parts of that network. Although there are some technology related meetups here, most focus on topics of broad appeal (*e.g. Speak English Dublin*) that are likely to attract members with disparate other interests.

4.2 Exploring Communities

The application of the OSLOM community finding algorithm yielded a model with 26 communities, ranging in size from 17 to 216 meetups. The mean community size was 65. Table 2 summarises details for the 15 largest communities identified in the normalised meetup co-membership network (full list is available in ¹). For each community, we report its size (*i.e.* number of meetups) and labels generated based on meetup name and description metadata, using the

Rank	Meetup Name	Score
1	Dublin Artificial Intelligence & Deep Learning	0.0659
2	Big Data Developers in Dublin	0.0658
3	Data Scientists Ireland	0.0654
4	Zalando Tech Events Dublin	0.0651
5	Machine Learning Dublin	0.0640
6	Data Science and Engineering Club	0.0639
7	Hackers and Founders Dublin	0.0638
8	GDG Dublin	0.0635
9	Dublin Startup Founder 101	0.0632
10	Dublin - Coder Forge	0.0630

(a) Weighted eigenvector centrality

Rank	Meetup Name	Score
1	Speak English Dublin	0.1053
2	AWS User Group Dublin	0.0467
3	New and Not So New In Dublin	0.0428
4	English Language Conversation with English Teacher	0.0349
5	Events, Drinks and Dancing in Dublin	0.0327
6	Python Ireland	0.0323
7	Dublin UX	0.0289
8	Dublin Indoor Football at The Soccer Dome	0.0260
9	Dublin LGBT Social Meetup	0.0229
10	Socializing Dublin	0.0218

(b) Weighted betweenness centrality

Table 1: Top 10 most central meetups in the Dublin meetup network, as ranked by (a) weighted eigenvector centrality and (b) weighted betweenness centrality.




Fig. 2: A visualisation of the meetup network for Dublin, where only intra-community edges are shown. Each node represents a meetup, where its size is proportional to the number of members in the meetup. Darker colour indicate a higher weighted degree of the nodes, numbers indicate community Ids.

approach described in Section 3.3. These labels indicate a diverse range of communities, covering topical areas such as technology, travel, self-help, music, and entrepreneurship. Despite the fact that the name and description fields typically differ considerably in length, it is interesting to note the relatively high level of overlap for the terms appearing on both lists for the same community.

To further explore the results produced by OSLOM, for each community we constructed a subgraph of the original network, and then ranked the nodes assigned to their community based on their centrality within that induced subgraph, as calculated by *weighted degree centrality*. The score for a node i is defined as the sum of the weights of the edges connecting i and its neighbours. Table 3 lists the top 3 most central meetups in each of the 15 largest communities.

These 26 subgraphs are shown in Fig. 2 generated using the MultiGravity Force Atlas 2 graph layout algorithm provided by the Gephi tool [4]. The connections between these subgraphs are due to the overlapping meetups between the communities, which makes the entire meetup.com community a single connected component. The size of each node in Fig. 2 is proportional to the number of members in that group; the colour of each node indicates its weighted degree.

The level of overlap between communities was not as high as might be expected. In total, 197 of the 1,482 meetups were assigned to more than one com-

Id	Size	Name Label	Description Label
5	216	hiking, international, wicklow, friends, yoga, book, culture, adventure, language, travel	fun, members, friends, time, hikes, free, social, friendly, looking, food
4	148	meditation, yoga, healing, spiritual, heart, sound, empowerment, soul, life, positive	healing, life, meditation, experience, self, energy, practice, spiritual, mind, mindfulness
7	137	data, user, science, tech, engineering, big, cloud, users, things, learning	data, programming, developers, community, code, science, software, technology, technologies, learn
17	118	user, tech, security, cloud, sharepoint, technology, game, software, data, crypto	data, learn, share, learning, developers, cloud, community, security, technology, software
14	84	business, digital, marketing, startup, entrepreneurs, network, job, professionals, innovation, market	business, marketing, digital, entrepreneurs, startup, market, network, owners, sales, job
22	80	yoga, meditation, workshop, stress, dun, laoghaire, camino, running, dance, therapy	yoga, life, body, meditation, class, health, classes, practice, energy, mind
3	78	startup, entrepreneurs, digital, lean, business, marketing, agile, growth, product, innovation	business, entrepreneurs, marketing, startup, networking, digital, lean, product, community, innovation
25	77	yoga, health, happiness, meditation, vegan, prayer, empowerment, circle, centre, self	yoga, life, meditation, help, support, healing, learn, world, health, work
10	71	user, mysql, traders, developers, tech, js, product, data, sprint, net	learn, product, developers, mysql, share, community, meetups, professionals, technologies, engineers
18	63	music, singles, rock, social, travel, south, international, fans, electronic, 30s	music, night, friends, fun, singles, rock, singing, love, members, sing
8	61	yoga, meditation, health, healing, classes, relaxation, self, body, light, sound	yoga, meditation, body, classes, life, mind, healing, health, practice, nature
21	54	empowerment, self, book, support, health, workshop, eating, therapy, life, development	life, world, diet, work, feel, learn, share, spiritual, ideas, find
15	53	circle, things, drinks, city, fun, hike, ladies, social, friends, book	drinks, friends, women, fun, book, food, wants, single, cinema, dinner
16	53	dance, dancing, yoga, classes, movement, salsa, fitness, class, set, handstand	dance, classes, dancing, fun, fitness, workout, 8pm, levels, class, movement
26	52	soul, prayer, network, life, healing, workshop, empowerment, biodanza, centre, body	life, god, healing, faith, spiritual, love, evening, work, reiki, chat

Table 2: Summary of the 15 largest communities identified in the Dublin meetup network. Community labels are generated based on the top 10 terms appearing in the meetup name and description fields.

munity. Of these, 176 appear in two communities, 20 in three communities, and a single meetup appears in four communities (*Headless Awareness Dublin*). Table 3 also reports the percentage of overlapping nodes in each of the 15 largest communities—*i.e.* how many nodes assigned to the community were also assigned to at least one other community. For some communities, a reasonably high proportion of nodes are overlapping (*e.g.* communities 21, 22, 25 in Table 3), while in other cases the vast majority of nodes belong exclusively to that community (*e.g.* communities 7, 10, 15 in Table 3).

To further illustrate the value of community finding, we present a more detailed analysis of a small subset of seven of the 26 communities found that have been manually identified as relating to technology. This subset is highlighted in




Fig. 3: The communities of technology focused meetups, each highlighted in a unique colour. Meetups in multiple communities are highlighted in orange.

Fig. 3, in which each community has a unique colour and meetups belonging to multiple communities are highlighted in orange. The community IDs from Tables 2 and 3 are also shown.

The biggest community in this subset (highlighted in red in Fig. 3) is community 7. *Data Science and Engineering Club* with 1,886 members has the highest weighted degree within this community, and shares more than 10% of its members with 19 other meetups in the community. The top five meetups with which the *Data Science and Engineering Club* meetup shares its members are: *Dublin Data Science* (30.8%), *Dublin Artificial Intelligence & Deep Learning* (22.1%), *Data Scientists Ireland* (22.0%), *Machine Learning Dublin* (21.0%) and *PyData Dublin* (19.1%). This shows a strong concentration around data science and machine learning. We see a similar concentration in community 2 (dark blue in Fig. 3) which contains meetups largely in the areas of cryptography and blockchain. The meetups in this community also share significant membership with community 17 (brown in Fig. 3).

It is interesting that this subset of the meetup community structure is connected to the remainder of the network largely by a small number of meetups. In particular *Entrepreneurs Anonymous Dublin*, *B.U.I.L.D business networking and more Dublin*, and *Dublin Ethical Entrepreneurs Applied Ethics Meetup* mediate the majority of the links of the technology communities to the rest of the Dublin meetup network. These nodes are also members of other communities.

Outside of technology there are other small, interesting, well defined social communities. For example, community 12 focuses broadly on art, the meetups in community 1 are largely focused on LGBTQ+ interests, community 9 contains meetups organised around playing board games, and community 6 is largely

Id	% Over	Top 3 Meetups
5	16.7	Kick Ass Adventures, Awesome Daytrips, Live Music Artistic and Cultural Events
4	12.2	Dublin Meditation Meetup, Dublin Self-Empowerment Coaching Meetup, Holistic Events Dublin
7	6.6	Data Science and Engineering Club, Dublin Node.js Meetup, Dublin Artificial Intelligence & Deep Learning
17	24.6	Social Software Development Meetup in Dublin, Amazic Ireland: Enterprise Docker and DevOps, Fairy Tale Artificial General Intelligence Solutions
14	32.1	Idea 2 Scale Dublin, SEO Strategy, Business Improvement & Innovation Network
22	35.0	Ki Yoga Massage, Dublin Wellness Meetup, Practical Ayurveda
3	34.6	Startup Grind Dublin, Dublin Startup Founder 101, Start-Ed Dublin: Free legal forum for Startups
25	35.1	Radical Forgiveness, Ki Yoga Massage, Dublin Grief Support Group
10	5.6	WebRTC Meetup Ireland, Serverless-Dublin, NetDevOps Dublin
18	23.8	Dublin Singles Party Meetup, Single Professionals Dublin, Dublin Nightlife
8	23.0	Holistic Health and Wellness, Gentle Yoga, Beginner's Meditation Course
21	42.6	Radical Forgiveness, Radical Human Expression, Dublin YouSA Meetup
15	5.7	Finding things to do in Dublin, 40 (or almost) and fabulous, Random things to do in Dublin
16	26.4	Stefano Argentine Tango, Samba Gafiera/Brazilian Zouk Dance Meetup, Dublin Dancing Meetup
26	25.0	Healingoptions4u, Radical Forgiveness, Pilgrims' Prayer

Table 3: The top 3 most central meetups in each of the 15 largest communities found. *% Over* shows the percentage of overlapping nodes in each community.

related to sports. The overlaps between these communities are also interesting. For example, community 12 has a relatively large overlap with community 1, and community 9 has a relatively large overlap with community 6.

5 Conclusions

This paper demonstrated the use of network analysis techniques to explore public data collected from the meetup.com platform, to characterise the fabric of a city. A co-membership network of meetups from Dublin, Ireland was constructed. This network allowed us to reveal the most important meetups in Dublin via measures of node centrality. To uncover the structure of the network at a higher level, community finding techniques were applied. By subsequently applying text analysis procedures to the aggregated metadata associated with each community, we have shown that thematically-coherent communities exist within Dublin's Meetup ecosphere. We also observed a limited degree of overlap between certain communities, where users might naturally share common interests.

Although the present study specifically focuses on Dublin's meetup network, using the same framework, the underlying communities of other cities could also be explored, and future work will develop a tool to support this. It would also be interesting to incorporate additional layers of meetup metadata into the network construction process. For example, some meetup groups are much more active than others, as some users are much more active than others, and this could be used to filter nodes in the network, and to influence the weight of edges.

References

1. Adamcsek, B., Palla, G., Farkas, I.J., Derényi, I., Vicsek, T.: Cfinder: locating cliques and overlapping modules in biological networks. *Bioinformatics* 22(8), 1021–1023 (2006)
2. Ahn, Y.Y., Bagrow, J.P., Lehmann, S.: Link communities reveal multiscale complexity in networks. *Nature* 466(7307), 761–764 (2010)
3. Amelio, A., Pizzuti, C.: Overlapping community discovery methods: a survey. In: *Social Networks: Analysis and Case Studies*, pp. 105–125. Springer (2014)
4. Bastian, M., Heymann, S., Jacomy, M.: Gephi: An open source software for exploring and manipulating networks. In: *Proc. ICWSM’09* (2009)
5. Bonacich, P.: Power and centrality: A family of measures. *American journal of sociology* 92(5), 1170–1182 (1987)
6. Brandes, U.: On variants of shortest-path betweenness centrality and their generic computation. *Social Networks* 30(2), 136–145 (2008)
7. Clauset, A., Newman, M.E., Moore, C.: Finding community structure in very large networks. *Physical review E* 70(6), 066111 (2004)
8. Edmonds, P.: Choosing the word most typical in context using a lexical co-occurrence network. In: *Proc. ACL-EACL’97*. pp. 507–509 (1997)
9. Fortunato, S.: Community detection in graphs. *Physics Reports* 486(3–5), 75–174 (2010)
10. Ghoshal, G., Barabási, A.L.: Ranking stability and super-stable nodes in complex networks. *Nature Communications* 2(1), 523 (2011)
11. Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proc. national academy of sciences* 99(12), 7821–7826 (2002)
12. Gmür, M.: Co-citation analysis and the search for invisible colleges: A methodological evaluation. *Scientometrics* 57(1), 27–57 (2003)
13. Greene, D., Sheridan, G., Smyth, B., Cunningham, P.: Aggregating content and network information to curate Twitter user lists. In: *Proc. 4th ACM RecSys workshop on Recommender systems and the social web*. pp. 29–36 (2012)
14. Jaccard, P.: The distribution of flora in the alpine zone. *New Phytologist* 11(2), 37–50 (1912)
15. Lancichinetti, A., Radicchi, F., Ramasco, J., Fortunato, S., Ben-Jacob, E.: Finding statistically significant communities in networks. *PLoS ONE* 6(4), e18961 (2011)
16. Lee, C., Reid, F., McDaid, A., Hurley, N.: Detecting highly overlapping community structure by greedy clique expansion. In: *Workshop on Social Network Mining and Analysis* (2010)
17. Newman, M.E.: Modularity and community structure in networks. *Proc. national academy of sciences* 103(23), 8577–8582 (2006)
18. Pizzuti, C.: Overlapped community detection in complex networks. In: *Proc. 11th Annual conference on Genetic and evolutionary computation*. pp. 859–866 (2009)
19. Small, H.: Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for information Science* 24(4), 265–269 (1973)
20. White, H.D., McCain, K.W.: Visualizing a discipline: An author co-citation analysis of information science, 1972–1995. *Journal of the American society for information science* 49(4), 327–355 (1998)
21. Xie, J., Szymanski, B.K., Liu, X.: SLPA: Uncovering overlapping communities in social networks via a speaker-listener interaction dynamic process. In: *Proc. IEEE 11th International Conference on Data Mining Workshops*. pp. 344–349 (2011)