

The New Blocs on the Block: Using Community Forums to Foster New Neighbourhoods

Elizabeth M. Daly, Dominik Dahlem
Smarter Cities Technology Centre
IBM Research - Ireland
{dalyeliz,ddahle}@ie.ibm.com

Daniele Quercia
Yahoo Labs
Barcelona, Spain
dquercia@acm.org

ABSTRACT

Research has consistently shown that online tools increase social capital. In the context of neighbourhoods Hampton and Wellman have shown that in newly developed areas residents effectively used mailing lists to connect with each other, circulate information, and ask for help. The research question of whether similar findings would hold in the larger context of a city for a long period of time is still open. To tackle this research question, we have gathered the complete dataset of the most popular neighbourhood online forum in Dublin. In this dataset, we have people sharing a common purpose (blocs) who live in the same neighbourhood and interact online to ask for help, engage in local activities, and, more generally, have a better understanding of their physical community. Our analysis highlights the particularly concentrated usage in newly established developments where a pre-existing community may be absent. Additionally, these communications provide a valuable resource to understand local issues relevant to the community.

Categories and Subject Descriptors

H.1.2 [User/Machine Systems]: Human factors

Keywords

Online Communities; Social Capital; Citizen Engagement

1. INTRODUCTION

Settling in a new neighbourhood often poses the challenge of getting acquainted with the surroundings and establishing a community feeling that provides a comfortable and safe place to live. People living in physical communities often turn to existing residents to learn about their area and neighbourhood. Wellman and Wortley found neighbourhood ties play an important role for people concerning specific support such as home improvements and assistance in emergencies [32]. However what happens when all neighbours are new? What happens if the sense of community has yet to be established?

A community that is better connected helps establish social norms, provides a source of information about local issues and enables

the residents to mobilise around community-specific issues. Online communities have emerged to ease the transition into a new neighbourhood and to strengthen the ties with the local communities. A great deal of discussion has emerged around the transformation of local community in the internet age, however much research suggests that online communication can play a key role in supporting local communities instead of supplanting the physical one [30, 24, 14]. From assistance in home repairs to private yoga classes and complaints about noise, people increasingly carry out elements of their lives online.

This paper focuses on data collected on a website designed to support neighbours build a community. Our main research contribution is to study whether and, if so, how the newly established neighbourhoods are linked with four main forum features that might increase social capital – activity, geography, topics, and emotions. We find that, compared to well-established residential areas, newly-developed ones...

- ... use the forum more often;
- ... connect with geographically closer neighbours;
- ... talk about their real-world needs (e.g., DIY);
- ... express themselves using positive and negative emotion words.

These findings suggest that this neighbourhood-based online forum has been a valuable source of social capital, particularly in the absence of an established physical one and highlight that there is much higher activity in newly established residential areas. We employ social network analysis metrics in order to understand the connectivity between community members. Understanding the different engagement levels across counties can assist policy makers wishing to use social media to improve communication with citizens and foster a sense of community. Topic analysis and sentiment analysis are used to gather evidence as to which issues are discussed by the members and identify the general feelings expressed online. As a result, the online forum can be a key tool for policy makers to have a finger on the pulse of the community, identifying the needs of residents such as requirements for schools and transportation infrastructure. In order to analyse the points above, we utilise the three main methodologies of social network analysis, topic analysis, and sentiment analysis to gain a deeper and principled understanding of the posts associated with residential areas.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '14, June 23–26, 2014, Bloomington, IN, USA.

Copyright 2014 ACM 978-1-4503-2622-3/14/06 ...\$15.00.

<http://dx.doi.org/10.1145/2615569.2615679>.

2. LITERATURE REVIEW

Social capital refers to features of social organisations such as networks, norms and trust that facilitate co-ordination, and co-operation for mutual benefit [22]. Online communities can provide a mechanism to foster social capital. In this section we review related work in the areas of online communities, the study of neighbourhoods and the study of technology and its potential role in neighbourhoods.

2.1 Online Communities

An online community can be defined as an Internet-connected collective of people who interact over time around a shared purpose, interest, or need [23]. Butler *et al.* examined what motivates users to invest in online communities focusing primarily on Listservs [2]. Though a large amount of social interaction on the web sustains pre-existing social ties between friends and family, a great deal of online interaction also occurs among people with no pre-existing relationships. They propose that people can benefit from online relationships as they create trust and provide a credible source of information, a form of social capital [22]. For example, people can now turn to online communities for important issues such as dealing with illness [5].

Online communities, however, only provides a technical infrastructure where social activity *may* take place [13]. As a result, it is important to understand what drives users to engage in a community. Ren, Kraut and Kiesler differentiate between identity and bond based online communities [23]. Identity-based communities are focused on the group identity as a whole where topics discussed are generally on-topic, and members feel a stronger association with the group rather than specific individuals. By contrast, bond-based communities tend to have more off-topic discussions and relationships are therefore more focused on ties between specific individuals.

Ellison *et al.* studied the use of Facebook in the context of a university campus. The study explored how it was used to build social capital, and was most useful to those with low self-esteem [7]. One could imagine a new neighbourhood community establishing itself could also benefit from online support networks.

2.2 The Study of Neighbourhoods

The study of neighbourhoods is not a recent phenomenon. Sociologists and urban planners have long studied the impact of social relationships, homogeneity, and group cohesion to understand the neighbourhood effect in the off-line world. Policy makers have begun to consider “local communities” and their interactions and evolution as a mechanism to tackle and gain insights into social issues such as crime, deprivation and social exclusion [8]. Lang and Hornburg examined how social capital affects housing and community development and found efforts to promote regeneration of disadvantaged neighbourhoods are more effective in areas where social capital is high [19]. Neighbourhoods play a strong role in promoting a sense of community and social cohesion [8]. Sampson *et al.* defined the term “collective efficacy” as a combination of social cohesion of a group and people’s ability to collectively intervene and act upon a common interest [25]. They found that collective efficacy is negatively associated with violence. Helliwell found that the ability to easily meet and interact with neighbours in friendly ways and a feeling of belonging to one’s community has an impact on trust in neighbours, which in turn has an impact on wellbeing [18].

2.3 The Study of Technology and Neighbourhoods

Hampton and Wellman ask “What is the Internet doing to local community?” [30]. Online communities and technology can be seen as a mechanism to foster social capital. Their research focused on studying a new housing development of Netville where many households were equipped with high speed internet access. Additionally, the community had its own mailing list where members circulated local information and asked for help. The authors found that the combination of online and off-line interactions enabled discussion around local issues and enhanced the sense of community. Resnick and Shah explored the impact of creating shared neighbourhood photo directories and found similarly positive results [24].

Hampton went on in 2007 to expand his study of communities in the e-Neighbourhoods project by examining four neighbourhoods (an apartment, a gated community and two suburban communities), three of which had mailing lists and the other had a neighbourhood website [14]. He found participation levels in online interventions varied depending on neighbourhood characteristics, especially depending on residential stability. The participation level of people living in the apartment block were low compared to the other communities. Through interviews the authors found agreement that there was not a strong existing sense of community. However, when queried residents expressed their desire to engage in additional social contact.

More recently, Harris and Flouch found that online communication through neighbourhood forums stimulates face-to-face connection, supports reciprocal assistance and favours [15, 16]. Both this study and Hampton’s found that lurkers do not benefit as much as active participants in the local online community and therefore understanding active participation in these networks plays an important role [15, 14, 14].

To sum up, previous work has not quantified the extent to which neighbourhood online forums are used by newcomers to “get started” in the context of an entire city. Our main research question is to study whether and, if so, how engagement with neighbourhood online forums in new developments differ from that in established areas. As opposed to current qualitative literature on the subject, we will conduct a quantitative study on how these forums are used differently across a variety of dimensions. To address this, we set out to study geographic-based communities that evolve around specific neighbourhoods in Dublin.

3. DATA DESCRIPTION

Neighbours.ie is an online community focusing around physical ones. The site started in 2006 as an Irish community based online forum organised around specific neighbourhoods and residential developments. The site had peak usage of over 3k posts a month in 2007 and has since declined somewhat in activity. Yet this data provides great insight into people’s concerns and how these are a reflection of the local urban setting.

3.1 Dublin Property Scene

In order to understand the importance of activity levels on forums surrounding one’s local neighbourhood, we need to shed some light around the economic climate of Ireland and Dublin, in particular in recent years. Financial innovation and liberalisation of the Irish market lead to a significant increase in available mortgage credit, even by international standards. The supply of credit had the consequence of pronounced activities in developing residential areas. As people moved into their new homes the sense of a local com-

munity was mainly absent. Online forums provided a means of coordinating on a multitude of issues that may arise with new buildings/developments and related infrastructural problems. This paper very much represents a social study to investigate concerns that are publicly discussed and to show how physical neighbourhoods and online communities augment each other.

3.2 Neighbours.ie

Our data consists of a crawl of the Neighbours.ie website for all County Dublin based forums from the start of the online forum in February 2006 until December 2011. This amounts to 20773 threads and 110317 posts contributed by 8801 authors¹. The forums are organised in a loose hierarchy of 20 root forums representing different post codes in Dublin and 5 forums related to the site in general. As a result, each forum can be linked to a geographic region and forums vary in terms of granularity. Some forums represent entire areas, others are sub-forums to reflect specific apartment blocks or developments. There are 718 unique forums, but only 474 of these have more than one thread which is typically the welcome thread created by the moderators. Generally, users posting on the respective forums are members of the site but they may also post as guests selecting a new user name every time. Over 98% of the posts are contributed by members.

We manually classified each neighbourhood into the following categories: area, street, mature development and new development. The date when a development became active was manually assigned with the help of data from newspapers articles, property websites and online forums. The neighbourhood categories are defined as follows:

- *Area*: An area is defined as a region where there is no singular residential development and the region goes beyond a single street or road. This includes areas defined by a posts code, town, suburb or village area. For example, Dublin 2 or South Docklands Area.
- *Street*: A street is defined as a forum where there is no clear association with a specific development and the title of the forum suggests it reflects a specific street. For example, Baggot Street or Mountjoy Square.
- *Mature Development*: A mature development is defined as a neighbourhood where the dwellings were built by a specific developer creating a collection of homes. This can include housing estates, gated communities and apartment blocks. For example, Wedgewood Estate or Pembroke Square Apartments.
- *New Development*: We define a new development as one that was built within the last 10 years, using the same criteria for a mature development where it refers to well defined collection of homes built at approximately the same time by the same developer. For example, Shelbourne Park Apartment or Bloomfield Park.

Though a large proportion of the forums were created by the website administrations, any user could create and add a new neighbourhood discussion forum. Figure 1 a) shows the break down based on the four different categories: 35% of the forums are centred around new developments, however mature developments are also present making up 25% of the forums. Therefore both types of neighbourhoods are relatively well captured on the basis of forum

¹For access to the data please contact the authors.

Distribution	Mean in mature	Mean in new	p-value
#posts	1.82	3.51	< 0.001
#posters	1.49	2.63	< 0.001
#threads	1.65	2.83	< 0.001
$\frac{\#posts}{\#threads}$	0.73	1.14	< 0.001

Table 1: Mann-Whitney test for the log of the means of our activity metrics: mature vs. new developments.

presence. The forums were manually geocoded and the author locations were inferred based on the weighted average of the location of forums they post to.

4. THE STUDY

We explore the relationship between online forums and the physical communities they represent. In particular, we now study how the presence of city newcomers is linked with four main forum features that might increase social capital – activity, geography, topics, and emotions.

4.1 New Neighbourhoods and Online Activity

Activity levels of an online community can be seen as indications of health of the community, how many people engaged in the forums and what the levels of engagement are. Prior work has shown that people turn to their social connections as sources of information and support [32]. Social capital researchers have found social ties with friends and neighbours are related to indices of psychological well-being, such as self-esteem and satisfaction with life [17]. Given new developments are less likely to have an established sense of community we propose that users turn to online communities in the absence of a physically established one.

We analyse the activity levels based on the type of community and the city regions. Figure 1 b) shows a boxplot of the number of posts in each forum for each neighbourhood type. The engagement levels for new developments is significantly higher than in mature developments with the least amount of engagement for forums that represent streets.

In order to further demonstrate statistical significance of variance in the numbers of posts, posters, and threads in mature developments and new developments we perform the Mann-Whitney test on the log of the means and obtain the results reported in Table 1. We confirm our hypothesis that the true differences in the means for all our activity metrics in mature developments are lower than those in new developments.

We find that forums associated with newly developed neighbourhoods tend to be more active than old neighbourhoods – they have higher number of posts, posters, and threads.

4.2 Newcomers and Their Social Network

Small densely connected networks are often associated with high levels of social and emotional support [22, 2]. In the context of neighbourhoods Sampson *et al.* defined “collective efficacy” as a combination of social cohesion of a group and people’s ability to collectively intervene and act upon a common interest [25]. In relation to neighbourhood forums this could range from coordinating around local issues such as the requirement for new schools to mobilizing and coordinating with management companies.

We investigate the topological relationship of users with respect to the forums they post to. The social network graph of the communities is generated to highlight clustered communities where en-

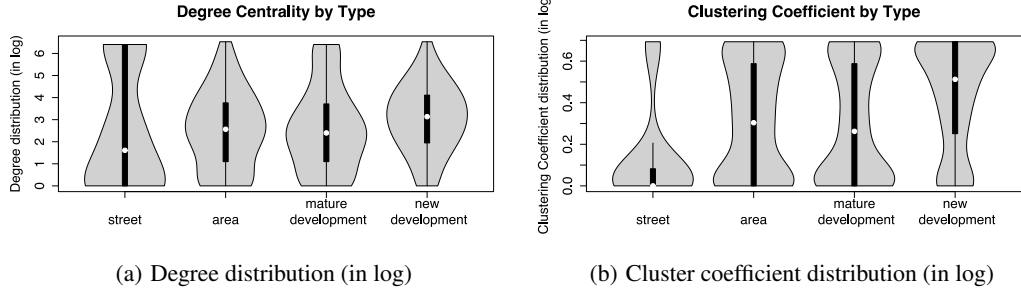


Figure 2: Network Statistics

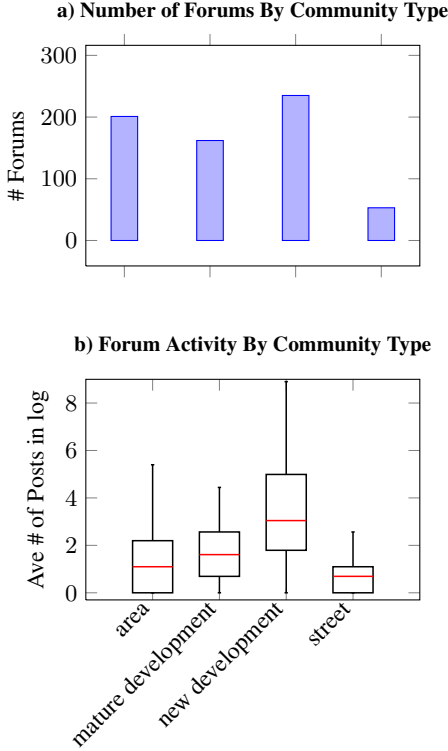


Figure 1: Forum Statistics

agement is strong. Given residents in new developments may need to organise themselves in order to coordinate with management companies and start to form a residential community we test the following hypothesis:

In order to explore the different topological social network across the different community types the online interactions are used to generate a social graph.

Each vertex in this graph represents a member if the website and users may post to multiple forums. A directed edge e_{ij} exists from user v_i to user v_j , if user v_i has posted to a thread to which also member v_j posted. More formally, this process of constructing the graph corresponds to the projection of the bipartite graph model of users and their relationship to existing threads onto the user-to-user graph. We also associate the number of overlapping threads between two users as the edge weight. Nearly 90% of the authors only posted to a single forum meaning their social links are only with those in the same neighbourhood forum. 4% of the users

showed no real affinity with a single forum and though it could be assumed these users are spammers, when examining these users we identified only 5 of these as members posting advertising services such as cleaning and painting. The remainder were outreach messages such as highlighting Fingal community services and surveys, seeking volunteers, youth club notices and education programs².

- *Degree Centrality* captures the total number of unique other members a person has engaged with by contributing to the same thread. The degree of a vertex v_i is the number of edges e_{ij} that are connected to the vertex [10].

$$C_D(i) = \sum_{k=1}^N a(i, k) \quad (1)$$

where $a(i, k) = 1$ if a direct link exists between v_i and v_k and $i \neq k$.

- *Node Clustering Coefficient* is a measure of how complete/connected the neighbourhood of a node is, where the neighbourhood of a vertex is defined as the set of vertices that are immediately adjacent to vertex v_i [29]. The clustering coefficient for v_i in an undirected graph is given by:

$$C_C(i) = \frac{2|e_{jk}|}{k_i(k_i - 1)} : v_j, v_k \in N_i, e_{jk} \in E \quad (2)$$

Figure 2 shows the distribution of degree centrality and node clustering coefficient. Forums representing new developments have the highest mean logarithm of the degree centrality however mature developments do show a number of members with a high degree centrality giving evidence to users contributing to a wide variety of threads. Forums representing streets have the highest variance with some highly connected users but a large number of users interacting with relatively few other members. However, figure 2 b) shows the clustering coefficient and cohesiveness of the street communities is very low. New developments on the other hand show the highest clustering of all the development types. It could be assumed that this is due to residents coming together on a small number of welcome threads, however, we found welcome threads tended to have very little engagement. Mature developments are bi-modal with a higher density around low and high clustering coefficients.

The Mann-Whitney test on the logarithm of degree and clustering coefficient shows a significant difference of the means for property types area and new development (p -value < 0.001), mature

²It should be noted these forums are moderated and therefore spamming posts have more than likely been removed meaning those that remain and are posted to many different forums are more likely to be genuine.

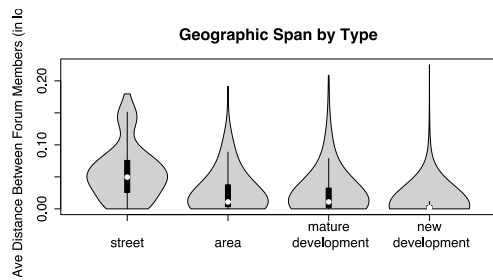


Figure 3: Distribution of average Spatial Distance Between Forum Members by Type

and new developments (p -value < 0.001), new development and street (only for the logarithm of the degree with a p -value < 0.001), and finally area and street (only for the logarithm of the clustering coefficient with a p -value < 0.001).

4.3 Newcomers and Geography of their Communication

Castells states that “people identify themselves primarily with their locality” [3]. Studies show that people communicate more through the Internet, and collaborate more effectively, when they are in closer proximity [31]. Liben-Nowell *et al.* supported this finding in the context of the popular blogging platform of LiveJournal [20]. Ren, Kraut and Kiesler differentiate between identity-based online communities where members feel a stronger association with the group and topic as a whole rather than focusing on individual connections in the case of a bond-based community [23]. Newcomers tend to be more welcomed in identity-based communities since the bond is between the community as a whole rather than specific members. Online communities based on neighbourhoods can be seen as identity-based communities where people socialise around local issues.

In order to explore how neighbourhood community types differ in their geographic property we examine the spatial distribution of the members actively engaged in the community. Each member is assigned a geographic location as a weighted average of the forum location the member posts to. Using this inferred location the average distance between posting members and the forum location may be used to calculate the geographic span of the forum. Figure 3 shows distribution of the average distance between the users who post to that community. New developments tend to have a lower geographic span compared to the other developments. Forums representing streets are associated with a relatively large geographic span showing more engagement from users not geographically clustered.

Population change and geographic span in new developments do not correlate, but the average geographic span of the posters for forums in new developments is *lower* than that for forums in old developments (Mann-Whitney test, p -value < 0.01) – so indeed “people identify themselves primarily with their locality”, which is more pronounced if those people are newcomers.

4.4 Newcomers and Topics

When studying the new housing development of Netville, Hampton and Wellman found that the neighbourhood mailing list enabled discussion around local issues [30]. More recently, the study commissioned by the Ford Foundation found similar patterns on community forums where discussions revolved around issues of local interest [9]. We propose that online neighbourhood communities are identity-based where members connect around local issues and

that members turn to the online community as a source of information and support around these issues.

Topic modelling is used to infer a topic associated with a set of words that co-occur significantly in a given text [26]. In our analysis we use Latent Dirichlet Allocation (LDA) to infer topics from our forum posts [1]³.

We aggregated all posts from a given thread into a single document and applied Porter’s Stemming Algorithm to the text. This text was then used to generate a topic model, selecting a topic number of 20. Table 2 shows a list of stemmed key words associated with each topic with a manually assigned label. As can be seen, the threads include residents discussing a variety of topics ranging from build related issues, management company discussions and community outreach posts. When looking at the key words associated with “Plumbing” it is interesting to note that along with words like “water”, “pump” and “boiler” is the word “anyon” which is the stemmed version of “anyone” which suggests this topic primarily arises in the form of a question asking “does anyone know how to...” or “does anyone know a good plumber?” supporting our premise that users are potentially turning to this online community to assist in problem solving or to find recommendations for these kinds of services when moving to a new area. We can also see evidence of local community posts around libraries, social tea events and education.

We focused our analysis on forums with a minimum of 20 posts, which yields 20567 threads to analyse. These threads were assigned a probability distribution over the individual topics. Each thread is then assigned the topic that is most represented in the thread. The topic probability distribution captures how strongly the posts in the thread reflect a given topic, some threads may possibly include a number of different topics and as a result some threads do not exhibit a strong relationship to one specific topic.

Figure 4 shows a heat map of the different thread topics for the different types of communities. Topics most prevalent to new developments are clustered to the left, while those most present in mature developments are clustered towards the right. The topic of plumbing, services, security issues and new build issues are the most discussed topics for new developments showing members are turning to the forum to gain information and potentially support around these issues. Services are an area of concern as they tend to be slow at coming online for new residential areas and members are seeking advice on service providers. Management companies are also actively discussed, as it is common practise for management companies to be hired to handle the administration of the new developments. Active topics for forums representing areas are community events and residents committee. Although forums representing areas are less active, those that are active are related to community engagement and organisation. Community events, home repairs and residential communities are active topics for mature developments as they maintain their more established neighbourhood.

4.5 Newcomers and Emotions

Members of neighbourhood forums share more than information only. They share norms and values within a culture of support and self help [28]. When connecting with companions in fortune and misfortune or in the organisation of bottom-up local activities, forum users connect with each other at an emotional level. This practically translates in the use of emotion words. Kivran-Swaine and Naaman found a link between the number of followers on twitter and the expression of emotions. They propose this could be due to

³We used the MALLET software to infer the topic models [21].

Topic ID	Label	Key Words
1	Plumbing	water, problem, heat, work, turn, pump, boiler, anyon, switch, time
2	Community Engagement	communiti, fingal, inform, pleas, group, librari, support, contact, local, partnership
3	Waste Management	bin, rubbish, clean, collect, estat, area, dump, green, day, garden
4	New Build	snag, week, block, move, anyon, told, builder, complet, hous, hope
5	Services	sky, dish, servic, ntl, broadband, phone, month, channel, connect, instal
6	Moderator/Welcome	forum, post, moder, ani, pleas, free, privat, send, neighbour, feel
7	Residents Committee	resid, meet, committe, issu, ani, post, pleas, contact, associ, peopl
8	Property Prices	price, year, properti, bank, market, interest, cours, hous, ireland, local
9	Reflective	peopl, don't, i'm, time, good, onli, live, realli, thing, ani
10	Family Area	kid, hous, dog, estat, car, area, road, live, children, peopl
11	Community Outreach	tea, event, fingal, contact, network, ticket, danc, citizen, pleas, parent
12	Community Events	school, club, educ, area, pleas, interest, class, gym, communiti, year
13	Security Issues	door, secur, block, night, car, alarm, gate, lock, nois, time
14	Management Company	manag, compani, pay, fee, year, charg, resid, servic, develop, agm
15	Volunteer	befriend, dublin, volunt, fundrais, servic, north, support, someone, experienc, train
16	Local	good, shop, open, great, food, veri, http, night, pub, nice
17	Home Repairs	floor, window, door, wall, hous, anyon, room, ani, paint, problem
18	Transport	road, bus, dublin, traffic, rout, servic, citi, stop, station, area
19	Parking	park, car, space, clamp, peopl, road, resid, visitor, spot, underground
20	Planning Permission	plan, develop, area, site, road, build, dublin, permiss, unit, propos

Table 2: Forum Thread Topic Keywords




Figure 4: Thread Topic Heat Map By Neighbourhood Category

shared emotional experiences creating more engaging content for people to consume [27]. Moving into a new area can be a combination of highs and lows and given many residents are experiencing this change at the same time.

In this section we assess the sentiment associated with a body of text. Global sentiment can be assessed and their temporal periods quantified using the micro-blogging platform Twitter [11]. Golder *et al.* used the Linguistic Inquiry and Word Count (LIWC) lexicon that provides measures of positive and negative affect. In our analysis we quantify sentiment according to [6] who measure societal-level happiness in Twitter streams according to numerical estimates of happiness of over 10,000 words. The happiness values associated with those words were collected through Amazon’s Mechanical Turk and does not rely on word stemming. The important aspect of their approach is two-fold. First, the word list is based on the most frequent words appearing in Twitter, Google Books, music lyrics, and the New York Times. Second, each word has been evaluated 50 times independently on a 9 point integer scale with respect to happiness where 5 represents a neutral value. Utilising this association of happiness to given words allows us to filter the neighbours.ie forum accordingly in order to assign a happiness value to each known word. The sentiment of each post on the forum is then independently quantified as the weighted average with respect to the frequency f_i of a word w_i given the estimated happiness value $h_{avg}(w_i)$ is set to:

$$h_{avg} = \frac{\sum_1^N h_{avg}(w_i)f_i}{\sum_i^N f_i}. \quad (3)$$

In the same fashion as in [6] we only consider words falling in the range of $5 - \Delta h_{avg} < h_{avg} < 5 + \Delta h_{avg}$, where $\Delta h_{avg} = 1.0$ to remove some high-frequency behaviour. Absolute measures of happiness as defined above, however, do not capture local differences in sentiment. To overcome this, we normalise the happiness score of a pool of texts by subtracting the global happiness score including all texts. This allows us to evaluate the sentiment of a particular forum with respect to the average overall sentiment of all texts. We analyse the sentiment on forums containing at least 20 posts, using the methodology discussed in [6]. It is important to note there were no forums focused around streets that contained more than 20 posts and therefore this category is omitted from this analysis. For the entire collection of posts, we calculated an average happiness measure based on the sentiment of the text and calculated an overall happiness value of 5.9 (where a value of 1 is completely negative and 9 is completely positive). We also compute a happiness value for each forum compared to the average. For a given forum, and a given pool of posts, we first find all posts from that forum and measure the average happiness of the subset of posts within the forum. The average happiness is then subtracted from that value. In this way, we see which forums diverge from the average.

There is little difference in the average sentiment across the different neighbourhood types, however, the true difference is seen in figure 5 which shows the distribution of thread post sentiment. New developments exhibit the widest range of emotions, fluctuating between negative and positive. This is partly due to the initial excitement of new residents and the expressed desire to form a positive new community contrasting with problems associated with new builds and management companies expressing frustration. A more positive sentiment is associated with developments focused around areas. This is mainly due to the prevalence of topics such as community engagement and community events being advertised as we have seen in our topic analysis.




Figure 5: Forum Normalised Happiness Distribution

The Mann-Whitney test shows a significant difference of the means of the sentiment shown in figure 5 between the property types of area and new development (p -value < 0.001) and mature and new developments (p -value < 0.01).

5. CASE STUDY: CELTIC HALL

During the Irish building boom of 2000s, a number of the property developments suffered building-related issues following completion. Our data set includes one that hit the media in October 2011. We use the forum of this development as a case study in evaluating the sentiment and topic analysis of the community and understanding what role the forum plays in connecting residents.

In August 2006, the forum was created and interestingly the development had not even been completed yet. The forum was used by people who had purchased property and were reaching out to their new community for assistance.

Hi, I was just wondering if anyone who bought in the first phase got follow up calls from the estate agent/builders.
- bigchicken. 22-08-2006

We notice that, the sentiment of the interactions for the first 4 months is positive in figure 6. Residents begin to discuss their communications with the management company and this topic is highly present in the first year. The forum also included threads related to the topic of community outreach aiming building a new local community.

SO whats gonna be the new local ??? I think we need a better boozer than the — inn ! me thinks the — residents bar. any better suggestions ? :-)
- namelock. 11-11-2006

In December 2006, there is a drop in sentiment because on December 12th a television show aired discussing building development and their lack of regulation. This resulted in much concern amongst the residents and also sparked the first discussion to mobilise the newly forming community.

Hi all, Perhaps it is time to organise the first Celtic Hall face to face Residents meeting and start making some real decisions on the issues so far experienced by those snagging, and the rather alarming issues raised by Prime Time last night.
- Harry. 12-12-2006

Not only did the members turn to the forum to coordinate, but they also expressed that the forum was a source of comfort.

hey all im feeling a little bit better then before i don't know why but earlier dismorning i was just in the height of it but this forum really helps....
-celtichall_audrey 12-12-2006

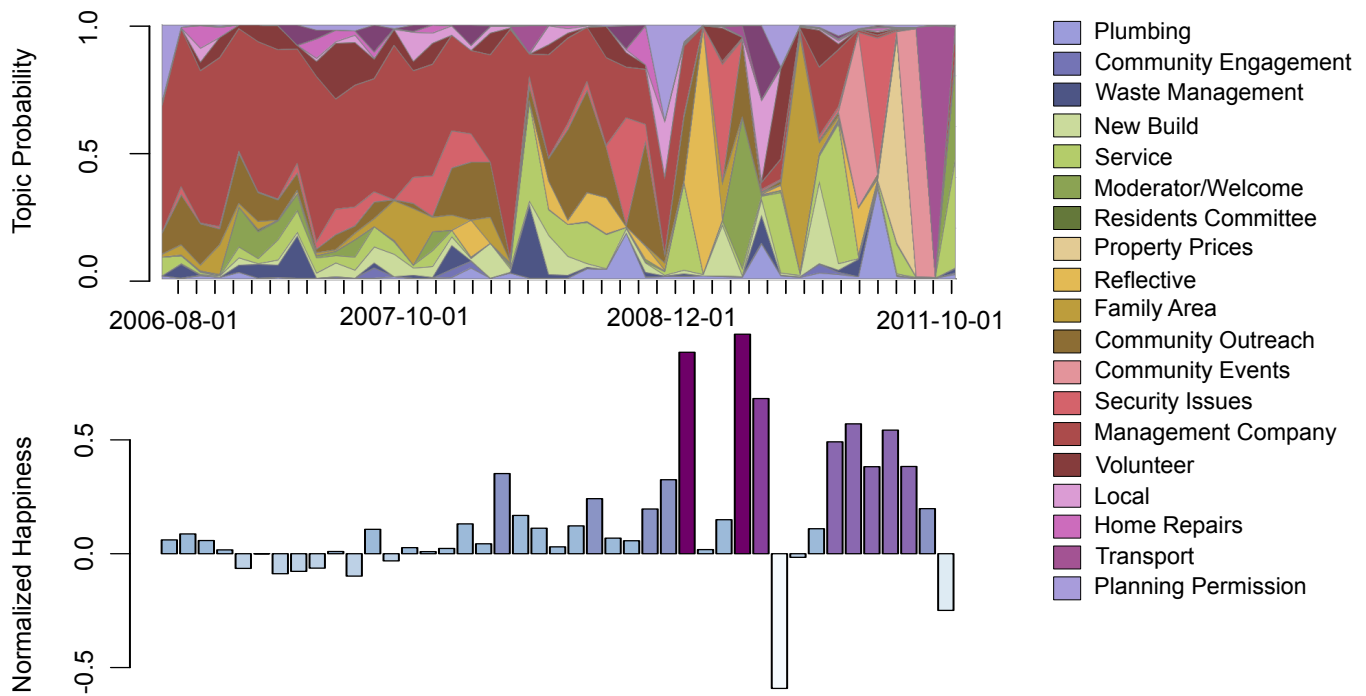


Figure 6: Celtic Hall - Sentiment versus Topic Distribution

The negative sentiment continues for the next few months as the completion date continues to be pushed back while residents use the forum to update each other with any new information. The sentiment returns to positive around May 2007 as residents start to take ownership and topics such as “new build”, “services” and “home repairs” start to be discussed and reducing the intense discussion on the topic of the “management company”.

Hey all Just back from my apartment for which I am now the key holder...I am absolutely delighted...To everybody who is still waiting to snag etc.. dont worry it will happen and soon we all will be neighbours. Best of Luck to you all and Im looking forward to meeting you soon !!!
-Macker. 23-04-2007

During this time, residents turn to the forum to ask questions about their apartment such as how to manage the heating or asking for advice on services. The community reduced in activity after residents had started to settle in. As can be seen the final month of December 2011 has a negative sentiment and it coincides with when residents had to be evacuated due to health and safety concerns. Though usage was less frequent following the move in, members returned to the forum to voice their concerns with the final postings in the dataset below:

I am posting a note regarding last weeks meeting with the Dublin City Council (15.09.11) and getting legal representation for owners as a joint effort, as this issue is of immense importance and one that will leave our properties unsell-able and unlivable but yet still mortgaged. We have less than 3 weeks to take action and create our own representation. Please reply to this if interested in meeting in the Hilton with a recommended solicitor who is willing to take this case. - R2D2 19-09-2011

As can be seen from this final posting, residents turned to the online community to share their concerns and offer members the chance to meet off-line and coordinate their actions. Though this case study provides primarily anecdotal evidence of forum use, it gives a sample of a number of real world interactions where the forum played an important role for neighbourhood members.

6. DISCUSSION AND CONCLUSION

Theoretical Implications. We have explored the online use of neighbourhood forums representing physical neighbourhoods. We have shown that even though the number of forums are relatively evenly represented by new and mature developments, activity levels and use are higher for new neighbourhoods. This finding supports the proposition that new residents may be turning to online forums in order to assist in establishing a community in the absence of a pre-existing physical one. Our findings suggest that given a digital forum of communication, residents of new developments are more likely to engage as an online community. Forums representing streets and areas are used significantly less. We propose this may be due to a lack of key topics that engage users such as interaction with management companies, shared experiences with build problems and fewer new residents arriving around a similar time frame. However, further investigation and analysis needs to be explored to support this hypothesis. Of the active forums representing areas, figure 4 shows the most popular topics included residents committees and community based discussions. This suggests a need to seed communities of more mature areas, streets and developments with topics and issues relevant to members in order to promote user engagement.

Practical Implications. Encouraging residents to interact through online forums that focus around local issues could have significant benefits for citizens and policy makers alike. Many social problems go beyond the control of government policy makers. However, enabling and supporting neighbourhoods to connect and foster a sense

of community can aid in shaping attitudes and social norms. Additionally, strong social networks have been linked to access to jobs [12], influencing health [4], and a source for help in times of need [32]. Online communities may also provide a valuable source of information when understanding the needs of citizens. For example, this work has highlighted the need for services and transportation in newly developed regions.

Limitations. This study has a number of limitations. The age profile of the members were not available to us and, given the preference for a younger generation to voice their concerns online, the results may disproportionately represent some citizens than others. Likewise internet access in the home may have levels of income bias. We also found a number of posts discouraging members to publicly discuss issues related to their community for fear of gaining a bad reputation and of negative effect on house prices. Use of the forums have declined somewhat in recent years. One potential avenue for investigation is to explore whether once the connection to the physical community has become more established and new issues arise less frequently, members may turn less to the online community. Further, we may find indications of reduced activity levels linked to a reduction in volatility of member sentiment over time. Another limitation is that our results do not speak to causality. Though we have found that engagement levels are higher for new developments we cannot necessarily conclude that these interactions lead to active support communities. One avenue for future work can include examining in more depth the text and content of member interactions seeking examples on supportive communication, which can also be followed up with surveys and interviews. Harris and Flouch found that online communication through neighbourhood forums did indeed stimulate face-to-face connection and support reciprocal assistance and favours [15, 16].

We have addressed the research question of understanding the online interactions of neighbourhood forums which work towards a ‘connected community’. For the first time, we have been able to conduct a *quantitative* study on an *entire* medium-sized city, and that has shown new insights on the geographic properties of these communities. Researchers have long been studying online communities in relation to physical neighborhoods, but the relationship between these two types of communities has been rarely investigated in a quantitative way (mainly because of lack of data). To answer the question asked by Hampton and Wellman of “what is the Internet doing to local community?” [30], we believe it is playing an important role in encouraging residents to come together around local issues, thus, building a new bloc.

7. ACKNOWLEDGMENTS

We wish to thank Gary Mc Ginty and Marco Capacchietti from the neighbours.ie team for technical assistance while crawling the site and providing permission to collect the data.

8. REFERENCES

- [1] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, Mar. 2003.
- [2] B. Butler, L. Sproull, S. Kiesler, and R. Kraut. Community Effort in Online Groups: Who Does the Work and Why? 2001.
- [3] M. Castells. *The power of identity : second edition with a new preface*. Wiley-Blackwell, 2010.
- [4] N. A. Christakis and J. H. Fowler. The Spread of Obesity in a Large Social Network over 32 Years. *N Engl J Med*, 357(4):370–379, July 2007.
- [5] K. P. Davison, J. W. Pennebaker, and S. S. Dickerson. Who talks? The social psychology of illness support groups. *The American psychologist*, 55(2):205–217, Feb. 2000.
- [6] P. S. Dodds, K. D. Harris, I. M. Kloumann, C. A. Bliss, and C. M. Danforth. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. Dec. 2011.
- [7] N. Ellison, C. Steinfield, and C. Lampe. Spatially Bounded Online Social Networks and Social Capital: The Role of Facebook. In *Annual Conference of the International Communication Association*, 2006.
- [8] R. Forrest and A. Kearns. Social Cohesion, Social Capital and the Neighbourhood. *Urban Studies*, 38(12):2125–2143, Nov. 2001.
- [9] T. F. Foundation. Inclusive Social Media Project: Participatory Evaluation. Technical report, 2010.
- [10] L. C. Freeman. Centrality in Social Networks Conceptual Clarification. *Social networks*, 1(3):215–239, 1978-1979.
- [11] S. A. Golder and M. W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, Sept. 2011.
- [12] M. S. Granovetter. The Strength of Weak Ties. *American Journal of Sociology*, 78(6):1360–1380, May 1973.
- [13] J. Hagel and A. G. Armstrong. *Net gain: expanding markets through virtual communities*. Harvard Business School Press, Boston, MA, USA, 1997.
- [14] K. N. Hampton. Neighborhoods in the Network Society the e-Neighbors study. *Information, Communication & Society*, 10(5):714–748, Oct. 2007.
- [15] K. Harris and H. Flouch. Social Capital and Cohesion. Technical report, Online neighbourhood networks study, 2010.
- [16] K. Harris and H. Flouch. The research context. Technical report, Online neighbourhood networks study, 2010.
- [17] J. F. Helliwell and R. D. Putnam. The social context of well-being. *PHILOSOPHICAL TRANSACTIONS-ROYAL SOCIETY OF LONDON SERIES B BIOLOGICAL SCIENCES*, pages 1435–1446, 2004.
- [18] J. F. Helliwell and S. Wang. Trust and Well-Being. *International Journal of Wellbeing*, 1(1):42–78, 2011.
- [19] R. E. Lang and S. P. Hornburg. What is social capital and why is it important to public policy? *Housing Policy Debate*, 9(1):1–16, Jan. 1998.
- [20] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(33):11623–11628, Aug. 2005.
- [21] A. K. McCallum. MALLETT: A Machine Learning for Language Toolkit., 2002.
- [22] R. D. Putnam. *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster, June 2000.
- [23] Y. Ren, R. Kraut, and S. Kiesler. Applying Common Identity and Bond Theory to Design of Online Communities. *Organization Studies*, 28(3):377–408, Mar. 2007.
- [24] P. Resnick and V. Shah. Photo Directories: A Tool for Organizing Sociability in Neighborhoods and Organizations. Sept. 2002.
- [25] R. J. Sampson, S. W. Raudenbush, and F. Earls. Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy. *Science*, 277(5328):918–924, Aug. 1997.

- [26] M. Steyvers and T. Griffiths. *Probabilistic Topic Models*. Lawrence Erlbaum Associates, 2007.
- [27] F. K. Swaine and M. Naaman. Network properties and social sharing of emotions in social awareness streams. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work, CSCW '11*, pages 379–382, New York, NY, USA, 2011. ACM.
- [28] M. Van den Boomen. Leven op het Net. De sociale betekenis van virtuele gemeenschappen. Technical report, 2000.
- [29] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, June 1998.
- [30] B. Wellman and C. A. Haythornthwaite, editors. *The Internet in Everyday Life*. Blackwell Publishers, 1 edition, Dec. 2002.
- [31] B. Wellman, B. Hogan, K. Berg, J. Boase, J.-A. Carrasco, R. Côté, J. Kayahara, T. L. M. Kennedy, and P. Tran. *Connected Lives: The Project*, chapter 8. Springer, 2006.
- [32] B. Wellman and S. Wortley. Different Strokes from Different Folks: Community Ties and Social Support. *The American Journal of Sociology*, 96(3):558–588, 1990.