

Google+ Communities as Plazas and Topic Boards

Michael J. Brzozowski
Google, Inc.
Mountain View, CA, USA
zozo@google.com

Phil Adams
Cornell University
Ithaca, NY, USA
pja22@cornell.edu

Ed H. Chi
Google, Inc.
Mountain View, CA, USA
edchi@google.com

ABSTRACT

Researchers have recently been focusing on understanding online communities in social networks that offer easy access to new audiences. How do online communities function within these social networks? In this work, we conducted a mixed-method study of public Google+ Communities and found two major types evident in both how users talk about them and how they appear to use them: *plazas* to meet new people, and *topic boards* to discuss common interests. This reflects two common motivations users cite in describing Communities: “meeting like minded people” and “finding great content”. We characterize these two types of Communities within Google+ using mixed methods including surveys, interviews, and quantitative analytics, and expose differences in user behaviors between them.

Author Keywords

Online communities; Online social networks; Google+

ACM Classification Keywords

H.5.3. Information Interfaces and Presentation (e.g. HCI): Group and Organization Interfaces

INTRODUCTION

For decades, online communities have provided a way for users to connect around their interests, ask questions, and provide support to one another. Before the Web, dial-up bulletin board systems and the WELL (Whole Earth 'Lectronic Link) community in 1985 [8] pioneered many community practices. Many forms of online communities flourished, including Usenet news, mailing lists, and web forums focused on specific topics such as health.

Recently, online communities have been coupled with online social networks (OSNs) such as Facebook, LinkedIn, and Google+, allowing users to interact with both their friends or contacts they follow, as well as fellow members of a community. This coupling enables us to examine the relationship and the differences between these two forms of social interaction. On one hand, communities enable users to circumvent the graph to reach others who share their interests. On the other, social network users may already be destined to connect to

others who share their cultural identity and values [12], so communities in an OSN may simply reflect existing ties.

Here we are interested in exploring how the combination with OSNs has changed the nature of online communities. We perform this work in the context of a relatively new social network, Google+¹, which is built around a directed social graph. In Google+, a user follows someone by adding them to one of his or her *Circles*; this action need not be reciprocated. The Google+ social graph is thus a hybrid of symmetric and asymmetric edges [24] just like Twitter, but with a higher clustering coefficient [13], reflecting a hybrid use to connect both with one's friends and with interests. In October 2013, Google announced the Google+ OSN has 300 million 30-day active users [7].

Google+ Communities² launched in December 2012, marketed as “for all the people you ought to know” [6]. Communities can be designated by their creator as *Public* or *Private*. Public Communities and their contents are visible to everyone, while Private Communities' contents are only shown to members. Additionally, the owner of a Community may elect to require new members to ask permission before joining. Users can find Communities from within Google+ by searching for topics, by scanning a list of system-selected “recommended” Communities, by receiving an invitation from someone they follow, or by following a link shared by someone they follow.

Our main research questions focus on the understanding of Communities within the Google+ OSN:

1. Why do people join Communities in Google+? Do their motivations align with prior studies of other online communities?
2. To what extent are Communities used as meeting places for new friendships to form? Are Google+ Communities used as a way to connect with new people?
3. Since the directed Google+ graph already affords following users' interests, how likely are people to be following fellow Community members?

In our study, we used a combination of several methods, including surveys, interviews, and quantitative analytics, to achieve some understanding of these questions. First, we coded 167 free-form responses from our survey to ascertain why respondents used Communities, and we interviewed 12 subjects in depth. Then we obtained a snapshot of the public membership rosters of 429,693 Public Communities

¹plus.google.com

²plus.google.com/communities

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

CHI 2015, April 18–23, 2015, Seoul, Republic of Korea.
ACM 978-1-4503-3145-6/15/04.

<http://dx.doi.org/10.1145/2702123.2702600>

in Google+ with more than two members, and analyzed the graph density and engagement metrics of those Communities.

We found most active users describe their motivations for joining Google+ Communities to be finding like-minded people and finding relevant content. Further analysis of Communities suggests these uses are supported by two main archetypes: *plazas* and *topic boards*. Many large public Google+ Communities fall somewhere on this spectrum. The findings suggest that successful communities can come in either flavor, but the prevalence of common interest forums suggest new ways of supporting these different types of communities in OSNs.

In the rest of this paper, we will first present a very brief overview of the rich history of online community research. Then we will outline the methods and results of the survey and interview studies. Next, we go into details of quantitative analyses of graph density and engagement metrics of a wide variety of communities. We further discuss the implication of these findings, and conclude by summarizing our findings.

RELATED WORK

An *online community* is a virtual community whose members interact with each other primarily via the Internet [25]. It's a simple information system where any member of the community can post content for other members to see and react to. Some communities restrict their membership, or have moderators with the ability to remove inappropriate content or people.

Many online forums have developed into communities of practice (COP) [29, 25], in which users are bounded by not just a common topical interest but also a common pursuit. Members in these COPs collaborate and help each other to improve their skills [25]. For example, members in many photography communities in Google+ often share tips, advice, and results of their work and failures with each other.

Understanding User Behaviors in Online Communities

Much has been written about building online communities—*e.g.*, how to bootstrap and maintain them, how to form social norms and regulate behavior, and how to encourage commitment and increase engagement. Kraut *et al.* offer a comprehensive synthesis of this research [10].

Sociologists distinguish two types of connections people have with the groups they join. In *common bond* groups, members feel more connected with each other than the group; in *common identity* groups, members feel more associated with the group identity than with other members [20]. An analysis of Flickr groups maps these categorizations into “social” and “topical” groups, respectively, finding that “social” groups have more comments per member [5].

Ridings and Green found that “information exchange” is the most common reason people join online communities, followed by either social support or friendship [22]. Wang *et al.* suggested that “informational support” is a key reason people joined a breast cancer support group, but “emotional support” from fellow members was a key reason they stayed [27].

OSNs and Communities

Online social networks (OSNs) developed much later than online communities, within the last 10–15 years. In OSNs, the focus is not on joining existing communities, but rather, finding specific friends or contacts to follow. Usually they feature a personalized stream of posts that does not focus on any one topic in particular, but is an aggregation of all posts made by contacts that the user follows. This leads to more sociality, but also much more complex social relationship structures.

For example, recent analysis of Twitter's follower graph suggests it exhibits properties of both a social and an information network [14], as users follow a mix of friends, celebrities, and news sources. To deal with this, Google+ users commonly engage in selective sharing to target posts at particular audiences who might be interested in them [9, 28].

Facebook Groups is one of the first combinations of online community features with an OSN. Park *et al.* found that Facebook Groups serve four primary needs: socializing, entertainment, status-seeking, and information [19]. Our notion of plazas and topic boards matches well with the two socializing and information uses of Facebook Groups. Other work has explored Facebook Groups being used for personal support (*e.g.*, with breast cancer) [1] and political advocacy [30].

In general, a common theme in online communities research is the categorization of online communities—indeed, Lazar and Preece contribute a clustering of classification schemes into four groups: (1) by community attributes (*e.g.*, social support); (2) by the software platform affording the community; (3) by connection to in-person communities (*e.g.*, a church group); and (4) by social boundedness [11]. Such classification is blurred by the evolution and overlapping of traditionally separate online spaces, *e.g.*, Groups and Communities being added to Facebook and Google+.

Following on this prior research, we wish to understand the blurring of the boundary between OSNs and online communities. We focus on understanding communities and their common use cases amongst Google+ users, and extending the categorization work to focus on two specific use cases that appear to be common in Google+ Communities. Do users see these communities as places to meet people, or to learn about a topic of interest? As relationships and interests are articulated by the graph in an OSN, we can also ask: Do users tend to join communities in which they socialize with other members they already know?

METHODS

We employed a mixed-methods approach since we are focused on understanding a wide variety of user behaviors at different granularities within Google+ Communities. Each successive method below refines the research questions we explore in earlier stages.

1. **Survey:** To understand how people perceive Google+ Communities, we deployed a short preliminary survey on Google+ and in key Communities. This provided a first broad understanding of how users perceive the Communities feature within Google+, including their initial use cases.

2. **Interviews:** Informed by the findings from the survey, we conducted a set of followup interviews, which provide for a closer look at why users behave the way they do in these Communities, what their intentions are, and how their goals manifest as actual behavior.
3. **Examining Google+ Communities:** To empirically measure the some motivations we saw in surveys and interviews, we analyzed activity data from Public Communities and the publicly visible portion of the Google+ follower graph. We identified outlier Communities to examine by hand. This provided a broad validation of trends that we found from the survey and interview studies.

SURVEY

In March 2013 we posted a short survey in English soliciting potential participants for future interviews. We reached respondents via snowball sampling: we posted on Google+ and in a few key Communities focused on discussing users' experience with Google+, and people reshared it on the network 110 times. We received 818 responses from 61 countries; the US, UK, Canada, India, and Germany were the most common, together covering 75% of respondents.

In addition to a few demographic questions, we asked those who confirmed they used Communities a single free-response question: "Briefly, what's your impression of Google+ Communities?" One of the authors then applied open coding to a randomly-selected 167 of these responses by hand.

Why People Use Communities

Fifty-nine of these responses were coded as describing why the respondent used Communities as follows:

Content

Forty (68%) of these respondents mentioned using Communities to find content around their interests. R160 said, "It's been very helpful to organize some of the things that I'm interested in."

R27 elaborated that it provides an alternative to using Circles:

Communities are great to find shared intetests [*sic*] without having to circle people. ... It also gives more room for discussions as people you are normally [not] connected with can find the post in a community and engage with you.

People

Twenty-seven (46%) of these respondents described the opportunity to "find like-minded people" or "meet people". While topical interests are often mentioned as a common trait used to find like-minded people, the responses in this category focused on actually finding or meeting *people*, rather than just getting interesting *content*. These are not typically random strangers, but people who share a common interest.

Some explicitly described the opportunity to *meet* new people. R21 said, "They provide a wonderful opportunity to meet and share with people having similar interests." R180 commented, "It further facilitates the process of meeting new people with similar interests, and condenses down to a forum where everyone can participate."

Others used vaguer terms like *find* and *connect*. R65 remarked, "I think this is a great tool to find people with related interests." "They seem like a good way to connect with people with similar interests", said R39.

In contrast to R27, R25 described using them as a way to find new people to follow: "Great way to find posts relating to a specific topic and people worth circling."

Relationship to the Circles Network

Another 27 respondents opined on the relationship of Communities to the broader Google+ OSN.³ This was somewhat controversial; ten people commented that they felt that discussion topics were now fragmented between their Circles graph and Communities. R120 described a pressure to join Communities to remain connected to content they were previously following:

...when they hijack conversations from the larger public community I don't like that. Say you have not joined every community that has to do with your main interest. Now the people who used to be in your stream have diverted their posts from the regular Public stream to a community. As far as I see, they're just gone. The person is still posting, and probably feels that their posts are still public (if it's a public community) but only those who have joined the community already will see the posts.

Meanwhile, another six commented that they preferred this separation.⁴ Google+ users employ selective sharing in part to avoid "spamming" uninterested people [9]. For example, R129 and R128 use Communities to target posts at specific audiences:

I like granularity to my interests so I can keep my specific nerdy things separate. My ukulele people don't want to hear about yarn or Ingress or writing or drawing. I like the separation. (R129)

I do love how the focus on certain aspects of one's interests can be segregated from one's main stream. I can get specific questions answered, discuss favorite topics, and not dilute my main stream with posts I don't want to share with everyone. (R128)

Summary

Two dominant motivations for Google+ Communities use emerged, covering 83% of those who described why they used Communities (49 respondents)⁵: a place for interesting content beyond the Circles graph (both to receive content and to post it), and to find or meet like-minded people; this is consistent with [22]. There still appears to be a large group of respondents who are more interested in meeting people than in merely getting access to interesting content.

³Other responses included: 88 describing general sentiment toward Communities, 30 citing specific frustrations, and 25 feature requests.

⁴The remainder of comments in this category related to details of the UI implementation and integration into Google+.

⁵Other uses mentioned included: 10 connecting an existing real-life group, 7 learning about a topic, and 6 having a place for discussion.

As we only had one author coding responses, and recruited through snowball sampling, it's not appropriate to infer the precise prevalence of these attitudes in the general population. However, we can still have reasonable confidence that a majority of respondents expressed one of these two uses⁶.

There's certainly a lot of overlap between these respondents—we coded eighteen (30%) as motivated by both content and people. Other, less prevalent, uses cited by respondents included learning and discussion, organizing an existing offline group, and connecting friends or family.

We explored these uses in more depth with interviews and data analysis.

INTERVIEWS

To get a sense of how people connected with each other in Communities, we conducted twelve one-hour interviews with Communities users in the United States who specifically reported having met someone in a Community. We sampled respondents recruited via both snowball sampling in key Google+ Communities and a broader user research panel. Participants were eight men and four women ranging in age from 18–60 from across the continental US.

Interviews were semi-structured, conducted by the same author, and video was reviewed by a second author. We derived themes from an *ad hoc* analysis of interview notes, agreed upon by two of the authors. Quotes were verified from interview videos.

Finding Communities

Participants articulated a variety of motivations for joining Communities, mirroring earlier findings from our survey. All expressed a desire to follow an interest and get quality content. Ten of the twelve specifically said they wanted to interact with other people around a shared interest.

Although they had joined as many as 50 Communities, none of these participants paid regular attention to more than five Communities. Of those five, they report typically feeling a strong connection to just one or two favorite Communities.

Some participants reported discovering Communities by seeking them out to indulge an interest. The Communities participants seemed to report the greatest sense of belonging to, however, tended to be those they received an invitation to from someone they admired or respected.

Orientation

First impressions can make or break a new member's engagement with online communities [2]. Some members were surprised to discover Communities that fostered collective participation [25]:

Joining a Community doesn't cost anything; you don't have to invest anything. So even if there were just interesting links and pictures, that'd be cool. I was looking just to see what was out there. But I found something more than that—something cool: at the end of each

⁶Even if up to 22% of responses were mis-coded, 36/59 would represent a majority of the sample ($p < .05$).

month everyone votes on a pattern and the next month everyone knits it. (P9)

Some Communities felt particularly tight-knit, supporting the use case of meeting new people. This quality is often readily apparent to newcomers:

I'm kind of surprised by the kind of and level of camaraderie that's immediately there... When new people come in people really latch onto them and find out about them—where do you live, what do you do, what's your lifestyle like. It really is like a family... there's just such support and interest. (P4)

Several participants reported that moderators played a large role in setting a tone for the Community. P12 said, "The mods are pretty good about making sure people play by the rules. You see the mods jumping in and letting people know." P5 added: "The photography communities all have really great mods, they're good about making sure people aren't posting junk and aren't posting other people's work."

When a "Group" Becomes a Community

Several participants articulated two distinct types of Communities: those that felt more intimate and private, and those that felt more like larger groups of people who merely subscribe to the same interest streams.

To me, the distinction between a private community and a public community is very clear. I like both, but you can tell the difference. My private⁷ one is smaller, they curate who comes in, and it feels like a community and I've made friends there. The big ones (Chromebook, for example) are active and I'll comment, but they are really like larger interest-based groups. (P10)

Another described the use of a Private Community as a place to seek support:

...and then there's my private community for really personal stuff. It's for eight people I've met through G+ where if I really need to talk about hard things, this is where I do it. Most of the communities I've been invited to are interest-based, topical. Not special purpose like this. (P11)

P12 was a member of a Community that organized group workouts over video Hangouts. The small-group interaction added a sense of belonging:

I thought it was pretty cool how [Jennifer] in one of the workouts was paying enough attention to the itty bitty little screens to say, "keep going [Michael]!". That makes it not just me sweating on my own in [Indiana], but part of a community.

Summary

Interviews revealed the importance of first impressions and moderator actions to set the tone for the types of interactions

⁷The Community described here is actually Public but members must ask to join it.

a Community will afford. Participants also perceived a difference between large “interest-based” groups and more intimate communities. Interestingly, the moderators’ engagements with the community foster a sense of intimacy and collectiveness.

EXAMINING GOOGLE+ COMMUNITIES

A significant portion of our survey respondents and interview subjects reported finding like-minded people in Communities, while others were driven more by the content. Likewise there appears to be disagreement about the extent to which people want Communities integrated with their social network. While many users *report* the desire to foster connections, we set out to *empirically* measure how likely people are to “know” others within their Communities.

To quantify this, we turn to the Communities themselves. We obtained snapshots of the member rosters of public Google+ Communities and public Circle edges⁸ in August 2013. Again, the Google+ social graph is directed: an *edge* represents one user following another by adding that user to a Circle, an action that need not be reciprocated⁹. All of this data is publicly visible to Google+ users. After discarding empty Communities, we found 429,693 communities with two or more members.

The size of Communities appears to follow a power law, with many small Communities and few very large ones. This is common in networks [16], and suggests many people feel empowered to experiment with creating their own Communities. Our informal observations suggest there are many tiny communities that reflect closer ties to a creator who was experimenting with starting a community. Thus, smaller communities likely reflect existing ties.

Dunbar proposed that humans can maintain at most 100–200 personal relationships, which has been reflected in the sizes of pre-industrial villages, businesses, and militaries throughout history [3]. We saw 17,896 Communities with more than 150 members, and a wide variety of Communities that exceed this by an order of magnitude: 2,449 Communities in our dataset had over 1500 members. Clearly, at these sizes one cannot be expected to know everyone in the Community.

Edge Density

To explore the extent to which people connect with each other in Communities, we examine the number of edges between members of the same Community. The more *dense* a Community’s graph is, the more likely members are to follow one another.

Only about 15% of the public Communities that *could* have any circle edges between members actually *do*, suggesting that Communities may be augmenting the social graph by providing access to people they would not be following otherwise and perhaps helping users break out of the “filter bubble” created by their networks [18].

⁸Google+ privacy settings allow users to make some edges private. Such edges are not included in this analysis.

⁹Reciprocal edges are not distinguished in this analysis; they’re simply counted as two following edges.

The maximum possible number of edges E in a Community with M members is $E_{max} = M(M - 1)$, a scenario where every member follows every other member. This theoretical ceiling is fairly elusive; no community with more than three members had more than $0.5E_{max}$ edges.

We fit a linear¹⁰ regression model to Communities with two or more members:

$$E \sim 1.812 \log M + .117M + (2 \times 10^{-7}) E_{max} - 1.651$$

The model’s R^2 value is 0.52, which is not a great fit, suggesting that the number of edges observed in a Community is not merely a function of its size, and some other hidden factors describe the likelihood of edge connection.




Figure 1: Number of internal edges in Public Communities v. number of members.

Figure 1 shows the relationship between members and edges. As communities scale up in membership, the number of edges increase on a log scale: that is, $\log E \sim \log M$. If the probability of any two people being connected with an edge is relatively constant, then we might expect E to correlate with the upper bound of edges possible, approximating $M(M - 1) \approx O(M^2)$. However, the number of edges does not keep pace. Therefore, in general, as Communities become larger, the likelihood of any two members being connected drops, reflecting our intuition about large communities being less personal.

The Outliers

Since Community edges are not merely a factor of their size, there must be at least one other factor that explains graph density. To explore this, we evaluated the Communities with the

¹⁰We also fit a log-linear regression model but it had a lower R^2 value.

largest residuals from this model—that is, the ones with surprisingly *many* or surprisingly *few* edges given the prediction from our model. Note that this naturally biases our analysis away from tiny Communities with few members. This suggests a spectrum model where the more (or less) residual edges a Community has, the more likely it is to behave like these extremes. Communities in the middle could exhibit characteristics of both, or neither.

We coded the top 30 on each side for the analysis that follows; for brevity only the top 10 on each side are listed in Tables 1 and 2. We found these extreme outlier Communities tended to exhibit common characteristics. Those with surprisingly many edges we call *Plazas*, while Communities with surprisingly few edges as *Topic Boards*. This is an admittedly arbitrary cutoff; we selected a sufficiently small group to so we could inspect them all by hand.

We heard in interviews that moderators play a large role in setting the tone for a Community. One of the key tools they have for defining their Communities and setting members' expectations is the "About this Community" descriptions, which prospective members can see when previewing a Community, and which remains readily accessible at the top of the Community's home page. We use these descriptions and an inspection of recent public activity in these Communities to describe how they function.

Plazas: Communities with Surprisingly Many Edges

Table 1 lists the Communities with the largest residuals from the model. Plazas fall into roughly two categories: meeting places and communities of practice.

Meeting places: Online communities can provide a sort of "virtual third place" for people seeking social interaction [26]. There appear to be a number of Communities providing this sort of "neutral ground" [17] where people who do not necessarily have a rich online social graph can meet others. Some communities make this explicit in its description:

Quero Ser Círculo [I Want To Be Circled] is a community that aims to group people to exchange ideas, make new friends and posts, is a great community for those who are new users of Google Plus. Because as everyone knows, when you're young, have few friends, and nobody shares or comments on your posts, stream completely still, it is very annoying. It's different here.

Common identity can be a powerful means of defining a community [21], and some plazas quite literally define themselves by shared identity. Eu Acredito em Deus! ("I believe in God!") invites members who self-identify as Christians: "JESUS this is your community, we created it because we love you and we are forever grateful for everything." Other Communities define themselves by national identity, like the simply-named Indonesia Community, which "is open to anyone who loves Indonesia".

Others, like Google+ Discuss, are targeted at new users looking to "chat about Google+" while meeting other newcomers. Meanwhile, others appear to target audience as wide as possible, like Earth, which invites users: "You live on Earth? Join!"

Communities of practice: A number of these dense Communities are also focused around the craft of photography and other arts (e.g., painting, sculpture). Many of them bear hallmarks of communities of practice [29, 25], serving as a virtual space for people wishing to practice the same craft. For example, Light Box explicitly encourages members to critique others' work if they wish to receive feedback on their own: "provide a short opinion/critique/vote on 15 eligible photos in the Voting pool...for each photo you submit". The moderators of Landscape Photography "highly encourage you to promote your own work and contribute to the knowledge of the other community members", while The Photo Community bills itself as "our neo-Montmartre".

There are also a few examples of communities of practice around online marketing, such as Hô +1 Google Plus ("+1 Google Plus Club"), a Community that specializes in search engine optimization and encourages members to share best practices.

These Communities often invite new members to introduce themselves, often through a dedicated category of posts, as seen in The Photo Community, HDR Photography, and Light Box.

For these communities of practice, one factor driving density may be the practice of sharing circles [4]; before Google+ added Communities, there were already informal communities of photographers following each other. It's possible that many of these edges existed before the formalized Communities did, and these Communities may be a reflection of existing clusters in the graph. This practice continues; TOP Photos on G+ has a category for members to share circles of photographers they like.

Topical networking: Of the top 30 Plazas, 26 are either *Meeting places* or *Communities of practice*. The remaining four are content-focused Communities that bear explicit calls to connect with other members. For example, Science on Google+ describes itself as

a community moderated by scientists, for all people interested in science, both professionals and the general public. The primary goal of this community is to bring real scientists to the public, for science outreach. A secondary and long-term goal is to create an environment that fosters interdisciplinary collaborations; thus, enabling and promoting cloud collaboration between scientists.

Topic Boards: Communities with Surprisingly Few Edges

Table 2 lists the ten Communities with the most negative residual edges—that is, those with significantly fewer edges than expected given our model. The 30 most negative Communities were all content-focused Communities we call *Topic Boards*. We draw this distinction in that these Communities make comparatively less effort to define their members by anything other than a topic of interest or to facilitate meeting people.

True devotees of *Harry Potter* or *Twilight* might self-identify as fans, and indeed, their respective Communities describe themselves as "for fans of Harry Potter on Google+" or "the

Community	Members	Edges	Residual	Type
Landscape Photography	164,645	77,515	52,810	Community of practice: photography
Google+ Discuss	77,034	40,925	30,706	Meeting place: new users
The Photo Community	41,464	32,614	27,401	Community of practice: photography
Hội +1 Google Plus (+1 Google Plus Club)	10,371	22,240	20,990	Community of practice: marketing
Hội Cần ++++++ (Need +++ Club)	9,883	20,924	19,733	Community of practice: marketing
HDR Photography	146,459	39,729	18,283	Community of practice: photography
Light Box	48,949	23,362	17,138	Community of practice: photography
Q.S.C - Quero Ser Circulado (I Want To Be Circled)	20,802	17,875	15,338	Meeting place: Brazil
TOP Photos on G+	4,838	14,892	14,308	Community of practice: photography
Photoshop and Lightroom Users	30,632	18,089	14,300	Community of practice: photography

Table 1: The ten Communities with the most positive residual edges (surprisingly many edges) according to our model.

Community	Members	Edges	Residual	Type
Harry Potter	173,564	2,326	-24,026	Interest: media
Gaming	120,247	2,289	-14,691	Interest: hobby
Forever Twilight	115,343	1,895	-14,280	Interest: media
Fitness & Weight Loss	94,599	2,048	-10,829	Interest: health
Joke of the Day	194,860	20,481	-9,932	Interest: humor
The Hunger Games	82,405	1,235	-9,783	Interest: media
Dogs	87,738	2,641	-9,183	Interest: hobby
Android	174,236	17,768	-8,709	Interest: technology
The Walking Dead Circle	71,053	1,470	-7,871	Interest: media
Cakes & Baking - The Cadbury Kitchen	58,245	886	-6,625	Interest: hobby

Table 2: The ten Communities with the most negative residual edges (surprisingly few edges) according to our model.

place where Twi-Hards can come”. However, they generally lack explicit calls to get to know other members or to interact with each other. It’s unclear whether an interest in *Harry Potter* is as deeply rooted a part of one’s identity as nationality or religion.

Topic boards promise rich streams of content focused on specific topics. For example, Gaming describes itself as “a bonanza of gaming news, reviews and discussion”. Joke of the Day offers “funny quotes, jokes, memes, photos, and good humor!”. This provides a way for members to find content around a topic of interest without much effort or expectation of engagement.

Topic boards also provide members with an audience for specific interests. Cakes & Baking invites members to “share and discover the most delicious recipes imaginable”. Dogs describes itself as “a place to share your dog pictures, stories, and experiences”.

Summary

In general, plazas tend to define themselves by declaring *who* belongs there, while topic boards tend to define *what content* belongs there. However, this is not a rigid binary categorization, but rather a spectrum, with some Communities bearing characteristics of both sides.

Also, while there are lots of photography communities of practice, topicality doesn’t necessarily determine a place on the spectrum. For example, while The Hunger Games is a topic board focused around the popular teen novels, the Chinese-language Community Google+ ShuPing (“Google+ Book Re-

view”) is a denser plaza of people invited to “share wonderful book reviews [and] recommend good books” .

User Engagement Analysis

Given these two types of Communities we believe to be qualitatively and measurably different, we next collected a set of metrics on per-user engagement in the 30 Communities with the most positive and 30 with the most negative residuals. For these 60 Communities, we measured and compared their total number of members, posts, and comments. None of these metrics are normally distributed, so we report the medians and apply a Mann-Whitney test to each pair. The data is summarized in Table 3.

Median of...	Topic boards	Plazas	<i>p</i>
Members	69,839	25,451	***
Edges	2,265	16,898	***
Posts	17,506	27,103	
Comments	49,706	39,092	
Posts/member	0.26	0.79	***
Comments/member	0.80	1.26	*
Comments/post	2.7	1.4	**

*** *p* < .001; ** *p* < .01; * *p* < .05

Table 3: Median values for basic metrics of the 30 topic board and plaza Communities we identified.

Membership size: In general, topic boards are larger than plazas (*p* < .001), as seen in Figure 2a. This makes sense: it’s hard to scale a plaza. But despite their larger size, Figure 2b shows topic boards tended to have fewer edges between




Figure 2: Distributions of members and edges for the topic boards and plazas in our sample




Figure 3: Posts and comments per member in topic boards and plazas in our sample

members ($p < .001$), reflecting the extremes of the spectrum we are sampling.

Activity per member: Both topic boards and plazas can generate lots of posts; the difference between them was not significant ($p = 0.4$). However, Figure 3a shows that plazas see significantly more posts per member ($p < .001$), suggesting that individual users are more likely to post in the Community, or more likely to post more. Likewise, plazas see significantly more comments per member ($p = 0.025$; see Figure 3b), consistent with findings in [5].




Figure 4: Comments per post in topic boards and plazas in our sample

Activity per post: Topic boards are hardly engagement wastelands, though; they tend to see more comments per post than plazas do ($p = 0.006$; see Figure 4). Therefore, while individual people may not contribute as often, for sufficiently large Communities *someone* is likely to see and comment on a post, even if it is a different *someone* each time. Sites like Digg have demonstrated this diffuse attention can effectively crowdsource evaluation of a high volume of content [31].

This pattern can help us partially distinguish between identity-defined plazas and identity-claiming topic boards. Table 4 shows four Communities—two plazas defined by religion and nationality, and two topic boards that describe themselves as for fans of specific novel series. The two plazas on the top have more posts per member and fewer members than the two topic boards on the bottom.

An analysis over the same engagement metrics revealed no significant differences between communities of practice and

	Posts/member	Members
<i>Plazas</i>		
Eu Acredito em Deus!	0.99	46,917
Indonesia	0.83	24,386
<i>Topic boards</i>		
Harry Potter	0.22	173,564
Forever Twilight	0.19	115,343

Table 4: Comparing plazas and topic boards claiming identity

meeting places. This agrees with survey and interview data in grouping these two Community styles together as plazas.

DISCUSSION

We presented a framework for understanding different forms of Google+ Communities. This division of Google+ Communities into plazas and topic boards parallels Park *et al.*'s work on uses and gratifications in Facebook Groups [19]. Their categories of *socializing* (e.g., peer support, meeting interesting people, talk about things with others, etc.) and *self-status seeking* (e.g., pressure to participate, make myself look cool) map well onto plazas, while *entertainment* and *information seeking* categories similarly fit topic boards. We extend this by showing these uses correspond to the prevalence or absence of edges observed within a Community.

Indeed, one possible explanation for this phenomenon may be that people are more likely to invite friends to plazas than topic boards. Each of these invites would increase the number of edges apparent inside the Community. While we did have anecdotal examples from interviews to suggest this, we did not have access to invitation logs to verify it.

Another way to understand the two styles of Google+ Communities is via the lens of *social awareness streams* [15]. Communities users driven by connecting with others may be similar to Naaman *et al.*'s "meformers", while those driven by content or an audience to share content with may be "in-formers". It is unclear the extent to which users can be strictly

separated into one of these two motivations, and it's possible the same users may be "informers" in some Communities and "meformers" in others.

Of course, this is not meant to be a strict dichotomy, but rather a spectrum. New Communities could become a densely-knit plaza or a popular topic board. Moderators help guide successful Communities toward their desired outcome.

Plazas reflect a scaled-up version of the desire to form connections with like-minded people articulated by our respondents in the survey and interviews. While greatest intimacy is afforded by smaller, more exclusive (and sometimes Private) Communities, and there may be a multitude of such dense Communities present, they don't appear in our list of extreme plazas because they're not large enough to have many residual edges.

At first glance it may seem that plazas are more compelling for users, especially given the potential for online communities to provide a social support net [1] or virtual "third place" [26], and the importance of common identity and interpersonal bonds in driving online community participation [21].

But not everyone wants to deeply connect with like-minded strangers online—some just want to get interesting content, or an audience to share their passion with. Topic boards are just as successful, since they provide a platform for people to get their content seen by others who share their interests, and they provide content that sparks interaction.

Obviously, deep personal ties can form even in ostensibly very topic-focused online communities [23]. Indeed, while users may enter a Community as a lurker with a topical interest, they may form bonds with other members as they begin to identify with the group. While we didn't have enough data to quantify this, we suspect that plazas may make this easier.

Implications

We see evidence that communities may provide an opportunity for people to form edges in an OSN. Some Communities explicitly encourage this behavior, and some appear to have unusually high graph density, suggesting that the Community itself may be a good source of contact recommendations.

However, people appear split on the extent to which their Communities should be integrated into their social graphs. Some prefer to keep their social network distinct from their interests, while others build bonds in Communities with the people they follow. This may suggest there is no "one-size-fits-all" approach to situating communities within an OSN, and that for some communities, nudges to follow fellow members outside the community may not be effective.

In any case, graph density is not necessary for lively activity within Communities. But it tends to afford a different type of activity: a smaller, more active core of members posting rather than a larger, diffuse crowd contributing more comments collectively on posts.

Finally, we propose that comparing graph density to a regression model based on its size is an easily implemented metric that could automatically detect whether a Community is

trending towards becoming a plaza or a topic board. This could provide guidance to moderators to help steer a Community's development, trigger adaptations in UI affordances, or suggest Communities for people who specifically want to meet others.

Limitations

There are some obvious limitations to our work, which was conducted on Google+, arguably a unique environment among OSNs. However, as mentioned above, our findings parallel existing work on Facebook [19] and Twitter [15], suggesting the patterns we observe may apply to other online community environments as well.

Our surveys and interviews were sourced through snowball sampling, which may bias responses somewhat towards more active users of Communities; lurker behavior is beyond the scope of this study. Also, they were conducted in English, which may limit global generalizability.

Another limitation is that the quantitative analysis considers only Public Google+ Communities on Google+, and in obtaining edge graph, membership, and activity data we are limited to only Public data. However, in interviews and surveys participants may choose to discuss Private Community activity. Since popular Public Communities are exposed in a variety of ways to potential members, Communities can form without existing edge ties. On the other hand, we expect Private Communities to behave differently, since typically users cannot see Private Communities unless they've been explicitly invited to them by a member.

Future Work

A natural extension is to evaluate the temporal ordering of edge creation and community joins to see in which Communities members actually followed each other after joining. This would paint a clearer picture of under what conditions Communities actually foster people meeting each other. A temporal analysis could also reveal whether Communities evolve into plazas or topic boards. Is it inevitable that sufficiently large plazas must become more topic board-like?

Another question is whether overall Community density is perhaps less important than membership in smaller, dense components within each community. Perhaps large topic boards contain a core component of tightly-knit members and a periphery of disconnected lurkers.

ACKNOWLEDGMENTS

We gratefully acknowledge Lichan Hong for technical advice. We thank Laura Granka and Dan Russell for helpful conversation, and our anonymous reviewers for their feedback and suggestions. We are also grateful to our survey respondents and interview subjects for their participation.

REFERENCES

1. Bender, J. L., Jimenez-Marroquin, M.-C., and Jadad, A. R. Seeking support on facebook: a content analysis of breast cancer groups. *Journal of medical Internet research* 13, 1 (2011).

2. Burke, M., Kraut, R., and Joyce, E. Membership claims and requests: Conversation-level newcomer socialization strategies in online groups. *Small group research* 41, 1 (2010), 4–40.
3. Dunbar, R. *How Many Friends Does One Person Need?: Dunbar's Number and Other Evolutionary Quirks*. Faber & Faber, 2010.
4. Fang, L., Fabrikant, A., and LeFevre, K. Look who I found: Understanding the effects of sharing curated friend groups. In *Proc. WebSci 2012*, 137–146.
5. Grabowicz, P. A., Aiello, L. M., Eguiluz, V. M., and Jaimés, A. Distinguishing topical and social groups based on common identity and bond theory. In *Proc. WSDM 2013*, 627–636.
6. Gundotra, V. Google+: Communities and photos, 2012. googleblog.blogspot.com/2012/12/google-communities-and-photos.html.
7. Gundotra, V. Google+ hangouts and photos: save some time, share your story, 2013. googleblog.blogspot.com/2013/10/google-hangouts-and-photos-save-some.html.
8. Hafner, K. *The well: A story of love, death & real life in the seminal online community*. Avalon Publishing Group, 2001.
9. Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. Talking in circles: selective sharing in Google+. In *Proc. CHI 2012*, 1065–1074.
10. Kraut, R. E., and Resnick, P. *Building successful online communities*. The MIT Press, 2012.
11. Lazar, J., and Preece, J. Classification schema for online communities. In *Proc. of the Association for Information Systems Americas Conference* (1998), 84–86.
12. Lazarsfeld, P. F., and Merton, R. K. Friendship as a social process: A substantive and methodological analysis. In *Freedom and control in modern society*, M. Berger and T. Abel, Eds. Van Nostrand, New York, 1954.
13. Magno, G., Comarela, G., Saez-Trumper, D., Cha, M., and Almeida, V. New kid on the block: exploring the Google+ social graph. In *Proc. IMC 2012*, 159–170.
14. Myers, S. A., Sharma, A., Gupta, P., and Lin, J. Information network or social network?: The structure of the twitter follow graph. In *Proc. WWW Companion* (2014), 493–498.
15. Naaman, M., Boase, J., and Lai, C.-H. Is it really about me?: Message content in social awareness streams. In *Proc. CSCW 2010*, 189–192.
16. Newman, M. E. The structure and function of complex networks. *SIAM review* 45, 2 (2003), 167–256.
17. Oldenburg, R. *The great good place: Cafes, coffee shops, bookstores, bars, hair salons, and other hangouts at the heart of a community*. Da Capo Press, 1997.
18. Pariser, E. *The filter bubble: How the new personalized Web is changing what we read and how we think*. Penguin, 2011.
19. Park, N., Kee, K. F., and Valenzuela, S. Being immersed in social networking environment: Facebook groups, uses and gratifications, and social outcomes. *CyberPsychology & Behavior* 12, 6 (2009), 729–733.
20. Prentice, D. A., Miller, D. T., and Lightdale, J. R. Asymmetries in attachments to groups and to their members: Distinguishing between common-identity and common-bond groups. *Personality and Social Psychology Bulletin* 20, 5 (1994), 484–493.
21. Ren, Y., Kraut, R., and Kiesler, S. Applying common identity and bond theory to design of online communities. *Organization studies* 28, 3 (2007), 377–408.
22. Ridings, C. M., and Gefen, D. Virtual community attraction: Why people hang out online. *Journal of Computer-Mediated Communication* 10, 1 (2004).
23. Rood, V., and Bruckman, A. Member behavior in company online communities. In *Proc. GROUP 2009*, 209–218.
24. Schiöberg, D., Schneider, F., Schiöberg, H., Schmid, S., Uhlig, S., and Feldmann, A. Tracing the birth of an OSN: social graph and profile analysis in Google+. In *Proc. WebSci 2012*, 265–274.
25. Shirky, C. *Here comes everybody: The power of organizing without organizations*. Penguin, 2008.
26. Soukup, C. Computer-mediated communication as a virtual third place: Building Oldenburg's great good places on the world wide web. *New Media & Society* 8, 3 (2006), 421–440.
27. Wang, Y.-C., Kraut, R., and Levine, J. M. To stay or leave?: The relationship of emotional and informational support to commitment in online health support groups. In *Proc. CSCW 2012*, 833–842.
28. Watson, J., Besmer, A., and Lipford, H. R. +Your circles: sharing behavior on Google+. In *Proc. SOUPS 2012*, 12:1–12:9.
29. Wenger, E. *Communities of practice: Learning, meaning, and identity*. Cambridge University Press, 1998.
30. Woolley, J. K., Limperos, A. M., and Oliver, M. B. The 2008 presidential election, 2.0: A content analysis of user-generated political facebook groups. *Mass Communication and Society* 13, 5 (2010), 631–652.
31. Wu, F., and Huberman, B. A. Novelty and collective attention. *PNAS* 104, 45 (2007), 17599–17601.