

Tracing Community Genealogy: How New Communities Emerge from the Old

Chenhao Tan

Department of Computer Science
University of Colorado Boulder
Boulder, CO, 80309
chenhao@chenhaot.com

Abstract

The process by which new communities emerge is a central research issue in the social sciences. While a growing body of research analyzes the formation of a single community by examining social networks between individuals, we introduce a novel community-centered perspective. We highlight the fact that the context in which a new community emerges contains numerous existing communities. We reveal the emerging process of communities by tracing their early members' previous community memberships.

Our testbed is Reddit, a website that consists of tens of thousands of user-created communities. We analyze a dataset that spans over a decade and includes the posting history of users on Reddit from its inception to April 2017. We first propose a computational framework for building genealogy graphs between communities. We present the first large-scale characterization of such genealogy graphs. Surprisingly, basic graph properties, such as the number of parents and max parent weight, converge quickly despite the fact that the number of communities increases rapidly over time. Furthermore, we investigate the connection between a community's origin and its future growth. Our results show that strong parent connections are associated with future community growth, confirming the importance of existing community structures in which a new community emerges. Finally, we turn to the individual level and examine the characteristics of early members. We find that a diverse portfolio across existing communities is the most important predictor for becoming an early member in a new community.

Introduction

“We all carry, inside us, people who came before us.”
— Liam Callanan

The tendency of individuals to flock together and form groups has led to continually emerging communities both online and offline. Websites that allow users to create communities at their own discretion (e.g., Facebook, Reddit, 4chan) provide a great opportunity to document this trend. For example, Fig. 1 shows a rapid increase in the number of communities ever since Reddit allowed users to self-organize into topic-based communities. Such growing trend poses an intriguing puzzle: where do these new communities come from?

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

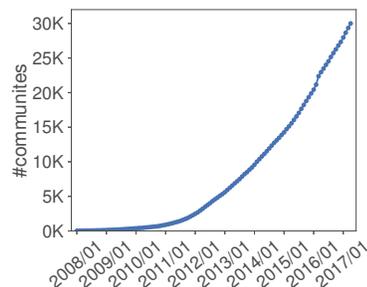


Figure 1: The total number of communities with more than 100 members on Reddit by each month: the number of communities has soared since 2008, when users started to be able to create their own communities.

In this work, we address this puzzle by considering each community as an entity, identifying the parents of a community, and building a genealogy of communities.¹ Although a battery of studies have leveraged online communities to investigate group formation and community growth (Backstrom et al. 2006; Kairam, Wang, and Leskovec 2012; Kossinets and Watts 2006; Liben-Nowell and Kleinberg 2008; Pavlopoulou et al. 2017), they tend to focus on a *single* community and seldom pay attention to the context that contains numerous existing communities. Our idea resonates with studies that explain the emergence of organizations by analyzing the synergy between a small number of close existing communities (Padgett and Powell 2012). For instance, Fleming et al. (2007) demonstrate how academic institutions and industry labs contribute to the emergence of high-tech companies in Silicon Valley and Boston. Our work aims to provide a computational framework for tracing the origin of a community among *all* existing communities.

Our main observation is that although every community starts from scratch (0 members), it does not emerge from a vacuum. In particular, members of a new community carry their own history, i.e., previous community memberships. The history of the early members allows us to understand where a community comes from and as a result, trace the origin of a community.

¹Anecdotally, GenWeekly has a series of articles on “genealogy of communities” (Edwards 2009).

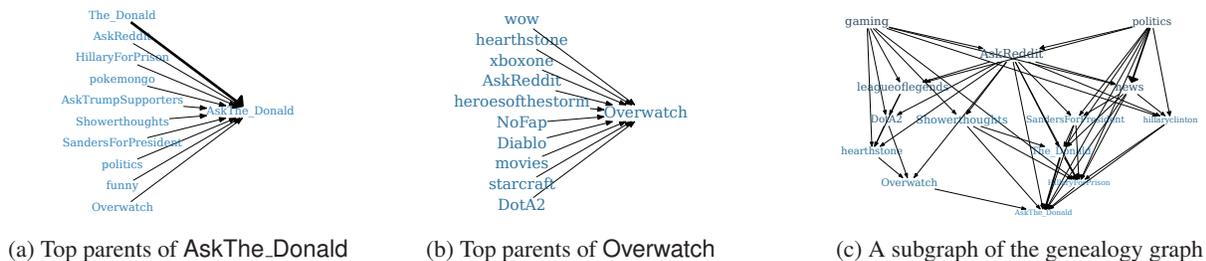


Figure 2: Genealogy graphs of example communities on Reddit based on the first 100 members. A directed edge indicates that there exist early members of the target node (“child” community) that were members of the source node (“parent” community); the thickness (weight) of an edge represents the fraction of such members. Node color represents the depth in the genealogy graph (darker color for older communities), while node size indicates *community size* measured by the number of members. Fig. 2a and Fig. 2b show the top 10 parents of AskThe_Donald and Overwatch respectively sorted by edge weights, while Fig. 2c presents the genealogy graph between a sample of communities starting from two of the first communities on Reddit (politics and gaming) to AskThe_Donald and Overwatch. To make the genealogy graph readable, we only present edges with weight greater than 0.01 (more than one members from the parent community).

Examples from Reddit. In order to build a genealogy of communities, we need the activity history of community members, which makes Reddit an ideal testbed. Fig. 2a and Fig. 2b show the top 10 parents of AskThe_Donald and Overwatch respectively, based on the first 100 members in each community. Not surprisingly, AskThe_Donald, a community where people ask Trump supporters questions, comes from The_Donald, a community for Trump supporters, and other politics related communities such as HillaryForPrison and SandersForPresident. Similarly, Overwatch, a community for a Blizzard game, comes from gaming communities such as wow, hearthstone, and Diablo. We also notice differences in edge weights between these two examples. Not any single parent of Overwatch is as dominant as The_Donald for AskThe_Donald.

Fig. 2c presents a broader view of the entire genealogy graph starting from the first Reddit communities, such as politics and gaming. We make several interesting observations. First, politics related communities and gaming related ones are clearly divided into two camps. Second, the connections are denser and weighted heavier for politics related communities, e.g., (The_Donald, HillaryForPrison). Third, despite being created more recently than politics and gaming, some communities such as AskReddit and Showerthoughts are among the largest communities on Reddit.

In this work, we provide the first large-scale characterization of such genealogy graphs. We further demonstrate the connection between the origin of a community and its future growth. We finally investigate the characteristics of early members to understand the basic element of our genealogy graphs.

Organization and highlights. We start by giving an overview of related work on group formation to provide further background for this work. We then introduce our dataset from Reddit, which spans over a decade.

We propose a framework for building genealogy graphs based on the first k members. Using this framework, we first track how a community *emerges* by examining the genealogy graphs when the number of members (k) grows from 0

to 100. We find that as k increases, the number of parents increases and the average weights of parents declines, indicating that the new community is becoming less dependent on any existing community. Second, we investigate how basic graph properties of the genealogy graph evolve as new communities emerge on Reddit. Intriguingly, despite the rapid increase in the number of communities over time, genealogy graph properties quickly converge to a stable state: the number of parents based on the first 100 members converges to ~ 180 and the dominant parent tends to contribute 10% of the first 100 members.

We further investigate how the origin of a community connects to its future growth. We find that genealogy graph information is useful for predicting how quickly community size grows (a 8.7% relative improvement in mean squared error). Our results show that strong parent connections are important for future community growth. This finding suggests that the emerging process of a community is analogous to complex contagion, e.g., the diffusion of political hashtags requires dense connections between early adopters (Centola and Macy 2007; Fink et al. 2015; Romero, Meeder, and Kleinberg 2011).

Finally, we formulate a prediction problem to better understand the characteristics of early members *at the individual level*. A diverse portfolio across existing communities turns out to be the most important predictor for becoming an early member, whereas community feedback and language use do not matter. Studies on early adopters of new products have found an overlap between early adopters, opinion leaders, and market mavens, who have information about many kinds of products and markets (Baumgarten 1975; Feick and Price 1987; Rogers 2003). Our finding suggests that early members of a new community tend to be market mavens instead of opinion leaders.

Related Work

Group formation and evolution has long been a focus in social science research (Lewin 1951; Coleman 1990). Here we discuss two most relevant strands of literature to provide

background for this work.

Group formation as a diffusion process. The process of group formation can be viewed as a diffusion process where joining a new group is analogous to adopting an innovation. A battery of studies have investigated the diffusion process of a single community or innovation (Aral, Muchnik, and Sundararajan 2009; Backstrom et al. 2006; Bakshy et al. 2012; Centola 2010; Centola and Macy 2007; Goel et al. 2015; Kossinets and Watts 2006; Liben-Nowell and Kleinberg 2008). The seminal work on group formation by Backstrom et al. (2006) shows that the likelihood of a person to join a group is associated with the number of her friends in that group; Centola and Macy (2007) propose the idea of complex contagion and suggest that the diffusion of certain behavior (e.g., political opinion) requires repeated contact and dense connections between early adopters.

Our work offers a new perspective in the context of constantly emerging communities: we view each community as an entity and focus on the relations between a new community and existing communities. In comparison, the basic unit of analysis from the diffusion perspective is the individual and the context in which a community emerges are existing networks between individuals.

Individual identity vs. community identity. Another closely related line of work is social identity (Tajfel 1982). One central hypothesis is that an individual’s self-perception derives from her group memberships. Thanks to the availability of datasets from online communities, researchers have become increasingly interested in studying user engagement in *multiple* communities as well as inter-group relations (Hamilton et al. 2017; Tan and Lee 2015; Vasilescu et al. 2014; Zhang et al. 2017; Zhu et al. 2014; Zhu, Kraut, and Kittur 2014; Hessel, Tan, and Lee 2016; Kumar et al. 2018). For instance, Tan and Lee (2015) construct the life trajectory of individuals using the communities to which individuals post on Reddit, and show that continual exploration is connected to users’ lifespan.

Instead of understanding individual identity from her life trajectory, our work attempts to examine the community identity by tracing where a community comes from. Community identity and individual identity are intertwined and their relation resonates with the two ecologies proposed in Astley (1985), population ecology and community ecology. We take the community ecology perspective and investigate the emergence of communities.

Finally, although our work focuses on the *emergence* of communities using explicit community structures, it is related to studies on online community building and loyalty (Hamilton et al. 2017; Hirschman 1970; Kim 2000; Kraut et al. 2012), and works on implicit community detection (Clauset, Newman, and Moore 2004; Girvan and Newman 2002; Yang and Leskovec 2013).

Reddit Communities

Our main dataset is drawn from Reddit, a popular website where users can submit, comment on, upvote, and downvote posts (Singer et al. 2014; Tan and Lee 2015). We refer to the difference between the number of upvotes and downvotes

as *feedback* for a post. We use posts from the inception of Reddit until April 2017.²

A brief history of communities on Reddit. Reddit started with a main discussion forum, reddit.com, and other subreddits such as science, politics, and gaming. In 2008, Reddit released a feature that allows users to create their own subreddits.³ Each subreddit focuses on a particular topic and has its own rules and norms, and thus functions as a community. Reddit now consists of tens of thousands of subreddits (henceforth *communities*), and the fact that we have access to the posting history of users enables our study.

This paper focuses on communities with more than 100 members so that we have a reasonable size of user base to trace the genealogy. As shown in Fig. 1, the number of communities has been increasing rapidly since 2008. There are over 30,000 communities with more than 100 members until January 2017. We consider a user as a member of a community if she made a post to that community.⁴ We refer to the number of users in a community as *community size*.

Our goal in this paper is to reveal the emergence of communities through building genealogy graphs between communities and exploring such genealogy graphs. The rapid increase in Fig. 1 is not even hindered by policies that raised the bar for creating communities by introducing additional user criteria in 2015.⁵ Such an ever-growing set of communities further motivates our research to understand the emergence of communities.

Characterizing Genealogy Graphs

In this section, we introduce the definition of genealogy graphs and explore their basic properties. We examine how genealogy graphs change as a community emerges, i.e., more users join the community. We also investigate the evolution of genealogy graph properties over time as users create more communities and find an intriguing convergence.

Building Genealogy Graphs

The key to building a genealogy of communities is to find a set of “parent” communities for each new community. We trace the parents of an emerging community by examining where its early members were active right before joining the

²This dataset is from <https://files.pushshift.io/reddit/>, thanks to J. Baumgartner. A small amount of data is missing due to scraping errors and other unknown reasons (Gaffney and Matias 2018). We checked the sensitivity of our results to missing posts with a dataset from J. Hessel; our results do not change after accounting for them. Details at <https://chenhaot.com/papers/community-genealogy.html>.

³<https://redditblog.com/2008/01/22/new-features/>.

⁴We only consider subreddits created until January 2017 so that even the newest subreddits have 3 months to accumulate members. Following Tan and Lee (2015), we only consider posting behavior and do not include commenting behavior because posting is initiated by users themselves, while commenting depends on other confounding factors such as the ranking system of Reddit.

⁵For more details, see https://www.reddit.com/r/help/comments/2yob6r/creating_a_subreddit/. This rule only limits who can create a subreddit and does not affect early members.

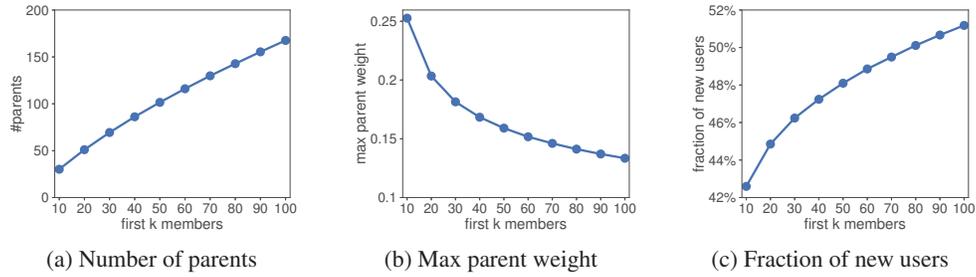


Figure 3: x-axis represents the number of members and y-axis represents basic properties of genealogy graphs. Error bars (tiny) represent standard errors. As the number of members increases, the number of parents of a new community increases in the genealogy graph, while max parent weight declines. Meanwhile, the fraction of new users, who do not have any recent community membership, increases. We also group communities by creation year and by community size to address concerns regarding Simpson’s Paradox and same trends hold in all conditions (Barbosa et al. 2016) (this is also true for Fig. 5).

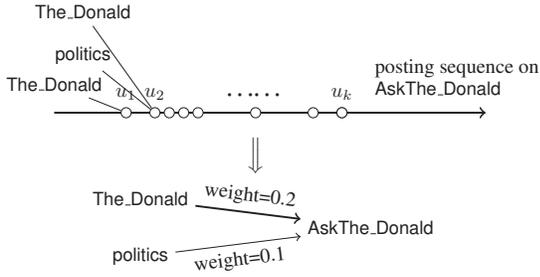


Figure 4: Illustration of the genealogy building algorithm for AskThe_Donald. The timeline shows the early members of AskThe_Donald and their posting sequence. Using the recent posting history of u_1 and u_2 , we can compute the weights for the genealogy graph at $k = 10$, assuming u_3, \dots, u_{10} did not make any posts in the previous month.

new community. In other words, we construct a community’s identity based on its early members’ recent community memberships. Specifically, we define parents of a new community j based on the posting history of its first k members in the month before they posted to community j . Fig. 4 presents an example for AskThe_Donald when $k = 10$. We define the weight on edge ij as the fraction of early members in j that were active in i :

$$w_{ij} = \frac{|\{u_r | 1 \leq r \leq k, u_r \text{ posted in } i \text{ recently}\}|}{|k|},$$

where i is the *parent* community, j is the *child* community, and u_r is the r^{th} member in community j . We identify the parents of all communities that were created after 2008, when users started to self-organize into communities. We approximate the creation time of a community by when the first post was made to a community. We only keep edge ij if community i was created before community j . To measure how dominant the top parent is (e.g., The_Donald for AskThe_Donald in Fig. 2a), we define *max parent weight* as $\max_i w_{ij}$ for child community j .

The number of early members (k) is an important parameter in this definition. We use small k values to focus on

the emergence of a community when the total number of members is small (≤ 100). For the same reason, we use an absolute count of k instead of a relative percentage. As k increases, a community’s identity takes shape. We track the emergence of a community by varying k from 10 to 100 and observe how genealogy graphs change.

Our approach is the first attempt to build genealogy graphs between communities. We will discuss future research directions in the concluding discussion. The extracted genealogical edges necessarily depend on our definition through previous community memberships of early members, so parent-child relations in our genealogy graphs indicate where the “child” comes from. It follows that “parent” and “child” do not have to relate semantically, e.g., Overwatch being a “parent” of AskThe_Donald in Fig. 2a.

From 0 Members to 100 Members

As the saying goes, “Rome wasn’t built in one day”, every community starts from 0 members no matter how many members it eventually has. We study the emerging process by tracking the genealogy graph from $k = 10$ to 100.

We study how basic properties of genealogy graphs between communities change in the emerging process of a community (Fig. 3). As the community matures and gains more members, additional members bring more past history and thus more parent communities in the genealogy graph. Meanwhile, the influence of parents decreases, indicated by the declining trend of max parent weight. The declining influence of parents is inherent to the construction of the child community’s identity.

Another important role that a new community plays is to attract new users who do not have any previous community membership.⁶ We observe an increasing fraction of new users as a community emerges: earlier members (small k) tend to come with past history, while a larger fraction of later members (larger k) come afresh.

⁶This includes both new users on Reddit and “old” Reddit users who did not post in the one month before posting in the new community, which can be considered reactivated “new” users.

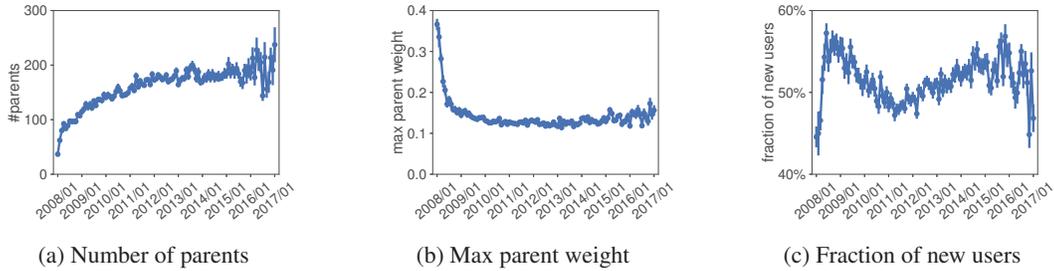


Figure 5: Evolution of genealogy graph properties over time. Each point represents the average corresponding properties of communities created in that month, and error bars represent standard errors. Both the number of parents and max parent weight converge rather quickly despite the rapid increase in the number of communities.

Dynamics over Time

In addition to changes during the emergence of a community, we investigate how genealogy graph properties evolve over time as the Reddit website as a whole grows. Despite the rapid increase in the number of communities (Fig. 1), graph properties converge to a stable state rather quickly. We focus our discussion on $k = 100$, but the observations are robust across choices of k .

The number of parents grows and converges to around 180 (Fig. 5a). As there are more communities over time and users continually explore new communities (Tan and Lee 2015), the number of parents increases because users carry more previous community memberships. However, the growth seems to have stopped since 2012. For communities created after 2012, the first 100 members tend to be previously active in around 180 existing communities.

Max parent weight declines and converges to 0.1 (Fig. 5b). As the number of communities increases on Reddit, early members of a new community tend to come from more diverse backgrounds. Therefore, max parent weight declines from ~ 0.4 to 0.1. However, this decline happens rather quickly. Starting from early 2009, max parent weight stabilizes at 0.1, suggesting that on average 10% of the first 100 members of a new community come from the same parent community. This observation holds despite that the number of existing communities grows from 100s to 10,000s since 2009.

The fraction of new users fluctuates over time (Fig. 5c). Unlike the number of parents and max parent weight, the fraction of new users presents a puzzling shape. The fraction of new users increased in early 2008 as Reddit just started to allow user-created subreddits. However, it went through a declining period until around 2011 and then started to grow again. One possible explanation for the increase in 2011 is the shutdown of the original main-reddit, reddit.com, which may reactivate some users to explore new communities, but this does not suffice to explain the growth in the fraction of new users from 2011 to 2014.

Predicting Community Growth

Having established the temporal dynamics of genealogy graphs between communities, we now investigate how the origin of a community relates to its future growth. We for-

mulate a prediction framework and study a community’s future growth after its number of members reaches 100. Our results consistently show that strong parent connections are associated with future community growth.

Problem Setup

In order to study how the origin of a community relates to its future growth, we develop two prediction tasks towards future community growth. First, inspired by Cheng et al. (2014) on predicting cascades, we use the median community size in our sample to characterize a community’s future growth and predict if a community’s future size is going to exceed the median community size (341 in our dataset).⁷ The advantage of this prediction setup is twofold: 1) it controls for the current size of a community and depends entirely on future community size; 2) it naturally leads to a balanced classification task. The majority baseline gives an accuracy of around 50%. We call this task *growth classification* for short.

The second task examines the rate of growth. We estimate how long it takes to reach the median for the subset of communities that exceed 341 members. We use $\log(t_{341} - t_{100})$ as the target variable and formulate a regression problem, where t_k indicates the first posting time of the k^{th} member. This regression task focuses on about half of the communities in the previous classification task. We refer to this task as *rate regression*.

For both tasks, in addition to using all the information from a community’s first 100 members, we constrain ourselves to using only the information from the first $k = 10, 20, \dots, 90$ members so that we can evaluate the marginal gain from additional members’ information.

Characterizing the Origin of a Community

Inspired by previous work on community growth (Ducheneaut et al. 2007; Kairam, Wang, and Leskovec 2012;

⁷Different from the observation in Cheng et al. (2014), although community size follows a heavy-tailed distribution, the hypothesis that it follows a power-law distribution is rejected when testing the goodness-of-fit (Clauset, Shalizi, and Newman 2009). We thus use the empirical median instead of an estimated value from the best power-law fit.

Zhu et al. 2014; Zhu, Kraut, and Kittur 2014), we consider the following hypotheses and features to characterize the origin of a community. Since both a large future size and a fast growth rate indicate community growth, a feature that is positively correlated with whether the future community size exceeds the median should be negatively correlated with growth rate. We first report significance testing results for single features when $k = 100$ (Table 1) and then examine the effectiveness of these features in prediction for $k = 10, 20, \dots, 100$.

Temporal features. We employ the community creation time and the average time gap between posts as features. We use the community creation time to account for the fact that the Reddit website has been growing. It is less likely for newer communities to exceed the median size on average (partly because of a shorter history), however, given that the community size exceeds the median size, the growth takes less time for newer communities. We expect smaller average time gap to be positively correlated with future community growth and this is confirmed in Table 1. These two features constitute a strong baseline in previous studies (Cheng et al. 2014; Kairam, Wang, and Leskovec 2012).

Basic parent properties. The next two feature sets derive from the genealogy graph. We first capture basic parent properties by the number of parents and edge weights. We also include the interaction of these two variables, i.e., the number of parents with weights at least 0.05 or 0.1. A large number of parents and heavy parent weights indicate strong parent connections, which may make a community more likely to grow. However, diffusion studies suggest that large-scale diffusion can happen with very shallow depth and involve many individuals’ independent adoption (Bakshy et al. 2011; Goel et al. 2015; Romero, Tan, and Ugander 2013). Our results show that max parent weight is associated with future community growth, both in growth classification and rate regression. The number of parents with substantial weights (≥ 0.1) is positively correlated in growth classification, whereas the number of parents is *negatively* correlated. These observations suggest that *strong parent connections help community growth but simply having many parents can sometimes hurt*.

Meta information of parents. We also dive deeper into the origin of a community and measure meta information of parents. We measure the meta information from two perspectives: size and language. We compute the average, min, max, and standard deviation of parent sizes using $\log(\text{number of members})$ in the one month before the new community is created. We expect larger parent sizes to associate with child community growth. Our results show that not all parents matter the same way. Average parent size is not significantly correlated in growth classification, but weighted average behaves as expected (larger weighted average parent size indicates larger likelihood to exceed the median community size and faster growth rate).

We measure pairwise distance between language models of the parent communities in the one month before the child community is created, using titles and texts in the posts.⁸ We

⁸We compute pairwise distance for only the top 20 weighted

feature	growth classification	rate regression
temporal information		
community creation time	↓↓↓↓	↓↓↓↓
average time gap	↓↓↓↓	↑↑↑↑
basic parent properties		
#parents	↓↓↓↓	—
#parents with weight ≥ 0.1	↑↑↑↑	↑↑
max parent weight	↑↑↑↑	↓↓↓↓
meta information of parents		
average parent size	—	↓↓↓↓
weighted average parent size	↑↑↑↑	↓↓↓↓
average language distance	↓↓↓↓	—
max language distance	↓↓↓↓	↓↓↓↓
new user		
fraction of new users	↓↓↓↓	—

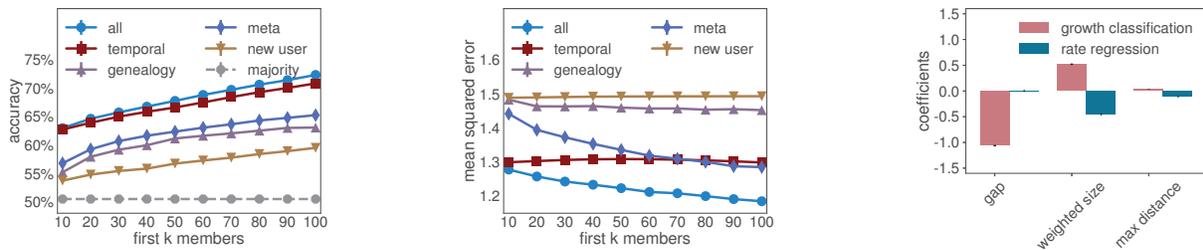
Table 1: Testing results for community growth. We show a subset of features for space reasons. \uparrow in growth classification and \downarrow in rate regression are positive signals for community growth. The number of arrows indicate p -values: $\uparrow\uparrow\uparrow$: $p < 0.0001$, $\uparrow\uparrow$: $p < 0.001$, \uparrow : $p < 0.01$, \uparrow : $p < 0.05$, the same for downward arrows; p refers to the p -value after the Bonferroni correction. t -test is used for growth classification, while Pearson correlation is used for rate regression.

compute the average, max, and standard deviation of pairwise distance and consider large language distance as a sign of diversity. We expect a diverse set of parents to be associated with community growth because Uzzi et al. (2013) show that atypical combination is related to scientific impact. However, this notion of diversity only seems effective for growth rate using max language distance. Large average language distance hurts the likelihood of exceeding the median community size and has no effect on growth rate. These results suggest that a closely related set of parents and maybe occasionally a small number of very different parents are associated with future community growth.

Fraction of new users. The fraction of new users is another way to capture the broadcast style diffusion (Bakshy et al. 2011; Goel et al. 2015; Romero, Tan, and Ugander 2013). We expect that the fraction of new users is positively correlated with future community growth. However, Table 1 shows the opposite. This observation further suggests that it is important for a new community to have strong roots in existing communities.

Summary. Overall, our feature testing results consistently suggest that having strong parent connections in the genealogy graph is positively correlated with future community growth. “Strong” parent connections can be reflected by a heavy edge weight, large weighted average size, small average language distance, and a small fraction of new users. The emerging process of a community is thus analogous to complex contagion (Centola and Macy 2007; Fink et al. 2015; Romero, Meeder, and Kleinberg 2011) and requires dense connections between parents.

parents for efficiency reasons. We also require that there are at least 100 unique members in that month to make sure that the language model is not too sparse.



(a) Growth classification performance.

(b) Rate regression performance.

(c) Coefficient comparisons.

Figure 6: Fig. 6a shows the prediction accuracy for different feature sets as the number of members grows (the higher the better), while Fig. 6b presents the mean squared error (MSE), and smaller MSE indicates better performance. Fig. 6c compares feature coefficients of these two tasks when $k = 100$ (features are standardized). Error bars represent standard errors.

Prediction Performance

We evaluate the predictive power of all features in both growth classification and rate regression. For each prediction experiment, we randomly split the data for training (70%), validation (10%), and testing (20%). We standardize each feature based on training data and use logistic regression with ℓ_2 -regularization. We grid search the best ℓ_2 coefficient based on performance on the validation set over $\{2^x\}$, where x ranges from -8 to 1 . We repeat this process 30 times and use Wilcoxon signed rank test to compare the performance of different feature sets (Wilcoxon 1945).

Performance in growth classification (Fig. 6a). Our classifier using all features outperforms the majority baseline and the performance improves quite significantly as k grows ($\sim 10\%$ absolute gain in accuracy by having the first 100 members vs. the first 10 members). The strongest single feature set is temporal information, which echoes the results in predicting cascades (Cheng et al. 2014). The performance gap between using all features and using temporal information increases as k grows, and the difference at $k = 100$ is statistically significant (72.3% vs. 70.8%, $p < 0.0001$). Meta information of parents is the second strongest feature set, whereas the fraction of new users performs the worst. This observation further confirms the importance of a community’s origin and suggests that the emerging process of a community depends less on broadcast-style diffusion in which many people join the group independently.

Performance in rate regression (Fig. 6b). Similar to the classification task, the performance improves as we get more information regarding early members (as k increases). However, different trends exist when comparing feature sets. Although temporal information remains a useful single feature set, its performance does not improve as k grows. There is a 8.7% relative gain in mean squared error for $k = 100$ by introducing features from the origin of a community (1.30 vs. 1.19). Our results suggest that the origin of a community is more effective in rate regression than in growth classification. One possible explanation is that the binary classification task is simple enough that temporal gap can already capture most of the signals.

Feature coefficients comparisons (Fig. 6c). In order to further understand the difference between growth classification and rate regression, we compare the coefficients in the lin-

ear classifier/regressor. The importance of average time gap drops significantly in the regression task (the coefficient’s absolute value drops from 1 to almost 0), but similar drop does not happen for features based on the origin of a community such as weighted average parent size and max language distance (as expected, their feature coefficients have different signs in these two tasks).

Who Becomes an Early Member?

We finally turn to the individual level and examine the characteristics of early members, because they are central to the emergence of a community. We formulate a prediction problem to differentiate early members of a new community from another user that was similarly active in a parent community.

Although early members have not been systematically studied as it is rare to have such activity sequence data that documents a community’s emerging process, many studies have looked at early adopters of new innovations and found an overlap between early adopters, opinion leaders, and market mavens (Abratt, Nel, and Nezer 1995; Baumgarten 1975; Feick and Price 1987; Rogers 2003; Ruvio and Shoham 2007). Market mavens have information about many kinds of products and markets, while opinion leaders provide information and leadership in specific products. Our results show that early members tend to be market mavens instead of opinion leaders.

Problem Setup

As active users may join a new community simply because of their high activity levels, we focus on understanding how factors other than activity levels relate to a user’s decision to become an early member at a new community. To do that, given a (parent, child) tuple in the genealogy graph and an early member of the child community (“positive”), we find a matching user with similar activity level⁹ in the parent community but did not become an early member of the child community (“negative”). The matching process leads to a balanced classification task, where the majority baseline gives an accuracy of 50%. We vary the number of early

⁹The nearest neighbor in #posts in the parent community in the month before the positive user joined the child community. We filter matches with distance greater than 5 (Stuart and Rubin 2008).

feature	significance	
	parent	global
#posts	—	↑↑↑↑
average time gap	↑↑↑↑	↓↓↓↓
feedback	—	—
language distance	—	↓↓↓↓
std language distance	—	↑↑↑↑
	interplay	
fraction of posts in parent	↓↓↓↓	N/A
community entropy	↑↑↑↑	N/A

Table 2: Testing results for predicting early members. Upward arrows indicate positive correlation with becoming an early member, while downward arrows suggest the other way around (↑↑↑↑: $p < 0.0001$, ↑↑↑: $p < 0.001$, ↑↑: $p < 0.01$, ↑: $p < 0.05$, the same for downward arrows; p refers to the p -value after the Bonferroni correction). Since behavioral features are used for the parent community alone and the entire Reddit website, they are respectively reported in the “parent” and “global” column.

members (k) from 10 to 100 to examine how early members differ as a new community emerges. We randomly sample 10K (parent, child) tuples, from which we collect 27.6K positive and negative users in total when $k = 100$.

Individual Characteristics

We design features based on an individual’s behavioral data both only in the parent community and in the entire Reddit website. Our features are derived from the following behavioral information:

- *Number of posts.* Although we try to control for #posts in the parent community, we expect more posts in Reddit to be positively correlated with becoming an early member.
- *Average time gap between posts.* More frequent posts indicate being more active and more likely to become an early member.
- *Community feedback.* Following Tan and Lee (2015), we measure community feedback by comparing a post’s feedback with the median feedback in that month to control for the differences across communities. Positive community feedback can be viewed as a sign of opinion leadership and may correlate with becoming early members.
- *Language distance from the community.* A user’s language distance from the community can relate to becoming an early member in either direction: a small language distance may suggest that the user fits in well and finds it hedonistic to explore, while a large language distance suggests that the user prefers something different and may move to a new community. We measure a user’s language distance from the community using cross entropy of her posts from the community unigram language model following Danescu-Niculescu-Mizil et al.; Tan and Lee (2013; 2015). In addition to average distance, we also compute the standard deviation of distances.

We first use hypothesis testing to evaluate single features when $k = 100$ (Table 2) and then experiment in a prediction framework to examine the effectiveness of these features.

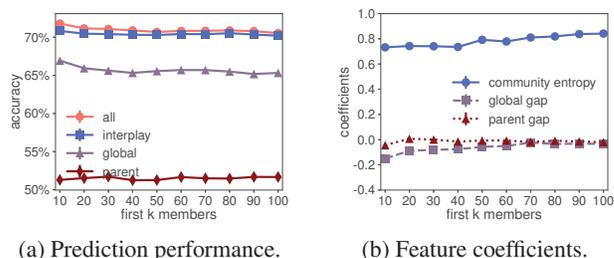


Figure 7: Predicting whether an existing user is going to become an early member in a child community. Fig. 7a shows the prediction accuracy for different feature sets as k increases. Interplay features dominate the performance. Fig. 7b shows coefficients of a feature in each feature set. Community entropy is more important than average time gap in the parent community and in the entire Reddit website. Error bars represent standard errors.

Small time gap in the parent community is negatively correlated with becoming an early member (the *parent* column in Table 2). Different from our hypothesis, our observation suggests that controlling for activity levels, users with a short burst of posts in the parent community are less likely to participate in the child community. Since we match users based on the number of posts in the parent community, it is expected that #posts in the parent community does not matter. However, little signal exists in language distance or community feedback in the parent community.

Users that fit well in language globally are more likely to become an early member (the *global* column in Table 2). Results on both the number of posts and average time gap are consistent with our hypothesis that active users are likely to become early members of the child community. Early members tend to fit well in language use in the entire Reddit website (smaller language distance). This is different from the characteristic of long-term staying users in Tan and Lee (2015) whose language use is more distant from the community than that of early-departing users. Intriguingly, larger standard deviation in language distance is correlated with becoming an early member, which suggests that diverse language use is a sign of early members. Again, community feedback does not seem to matter, which suggests that early members do not tend to be opinion leaders.

Early members have a diverse portfolio across communities (the *interplay* column in Table 2). We use interplay features to capture a user’s behavior in the parent community in the context of the entire Reddit. We compute the fraction of posts in the parent community and the entropy of posts across existing communities. We find that a diverse portfolio across existing communities is associated with a large likelihood to become an early member, indicating that early members share properties with market mavens.

Prediction Performance

We use logistic regression with ℓ_2 regularization and the same prediction setup regarding data split, feature normalization, and grid search as in predicting community growth.

Interplay features dominate (Fig. 7a). It turns out that although there are only two features in interplay features, they dominate the performance in predicting early members regardless of choices of k . The performance of using all features is relatively *stable* across choices of k , which suggests that characteristics of early members are insensitive to choices of k . Using all features slightly outperforms using only interplay features when $k = 100$ (70.5% vs. 70.2%, $p < 0.0001$). Features based on the parent community provide the worst performance, barely above the baseline, indicating that community feedback and language use in the parent community do not provide much signal when we control for user activity.

Feature coefficients (Fig. 7b) show a similar story that early members of a community are similar to “market mavens” in the domain of innovation adoption, who have information about many products and markets (Feick and Price 1987). The coefficient of community entropy dominates behavioral features based on the parent community alone and the entire Reddit website. The coefficient of average time gap in the parent community, for instance, is very close to 0, which explains the low performance of “parent” in Fig. 7a.

Our results show that although our task is set up based on a (parent, child) tuple in the genealogy graph, information regarding the user as a whole is valuable for predicting early members. Therefore, it is important that we take a holistic view of a user’s past history and consider the complete sequence of user behavior if possible.

Concluding Discussion

In this work, we study the emergence of communities and present the first large-scale characterization of genealogy graphs between communities, which are defined using previous community memberships of a community’s early members. We show that as a community emerges, the number of its parents increases and parent weights decrease. Intriguingly, despite the fact that the number of communities on Reddit increases rapidly over time, the number of parents and max parent weight converge to a stable state rather quickly. We also demonstrate the effectiveness of using a community’s origin to predict its future growth, especially in predicting growth rate. We find that strong connections with parent communities are related to a community’s future growth, which echoes the idea of complex contagion that the diffusion of certain behavior requires dense connections between early adopters (Centola and Macy 2007). Finally, we explore the characteristics of early members and find that a diverse portfolio across communities is a key factor.

Our work constitutes a first step towards understanding the emerging process of a community in the context of existing communities. Our observations open up future research directions. For instance, the convergence of basic graph properties over time may relate to the cognitive limit of individuals (Dunbar 1992) or the nature of building a new community, namely, the emergence of a new community requires a particular structure among early members (~10% of early members share experience in an existing community). It remains an open question to explain such convergence and

investigate whether similar convergence happens on other platforms or based on other membership definitions.

Furthermore, the notion of genealogy graphs has broad implications for understanding communities and requires much more exploration beyond our proposed approach. Although we focus on a small constant number of early members to examine the emerging process, the emerging process is inherently dynamic and requires varying user bases and time scales across communities. Future research may develop a robust way to adapt the definition of early members to different communities and even a general method to identify different stages of a community including the emerging process. Another interesting question is to examine the differences in a community’s origin and the characteristics of early members across different types of communities, e.g., communities on different topics (politics vs. gaming), or common-bond groups vs. common-identity groups (Prentice, Miller, and Lightdale 1994; Ren, Kraut, and Kiesler 2007). Last but not least, each member may not contribute equally in a genealogy graph. The nature of an edge between a parent community and a child community can further differ depending on why these members join the child community. It is a promising direction to build genealogy graphs with rich information on edges.

Acknowledgments. We thank J. Hessel, B. Keegan, V. Lai, L. Lee, N. Li, M. Naaman, A. Sharma, S. Zhang, and anonymous reviewers for helpful comments and discussions. We thank J. Baumgartner and J. Hessel for sharing the dataset that enabled this research. This research was supported in part by a gift from Facebook.

References

- Abratt, R.; Nel, D.; and Nezer, C. 1995. Role of the Market Maven in Retailing: A General Marketplace Influencer. *Journal of Business and Psychology* 10(1):31–55.
- Aral, S.; Muchnik, L.; and Sundararajan, A. 2009. Distinguishing Influence-based Contagion from Homophily-driven Diffusion in Dynamic Networks. *Proceedings of the National Academy of Sciences* 106(51):21544–21549.
- Astley, W. G. 1985. The Two Ecologies: Population and Community Perspectives on Organizational Evolution. *Administrative Science Quarterly* 30(2):224–241.
- Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group Formation in Large Social Networks: Membership, Growth, and Evolution. In *KDD*.
- Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone’s an Influencer: Quantifying Influence on Twitter. In *WSDM*.
- Bakshy, E.; Rosenn, I.; Marlow, C.; and Adamic, L. 2012. The Role of Social Networks in Information Diffusion. In *WWW*.
- Barbosa, S.; Cosley, D.; Sharma, A.; and Cesar Jr, R. M. 2016. Averaging gone wrong: Using time-aware analyses to better understand behavior. In *WWW*.
- Baumgarten, S. A. 1975. The Innovative Communicator in the Diffusion Process. *Journal of Marketing Research* 12(1):12–18.
- Centola, D., and Macy, M. 2007. Complex Contagions and the Weakness of Long Ties. *American Journal of Sociology* 113(3):702–734.

- Centola, D. 2010. The Spread of Behavior in an Online Social Network Experiment. *Science* 329(5996):1194–1197.
- Cheng, J.; Adamic, L. A.; Dow, P. A.; Kleinberg, J.; and Leskovec, J. 2014. Can Cascades be Predicted? In *WWW*.
- Clauset, A.; Newman, M. E. J.; and Moore, C. 2004. Finding Community Structure in Very Large Networks. *Physical Review E* 57(066111).
- Clauset, A.; Shalizi, C. R.; and Newman, M. E. 2009. Power-law distributions in empirical data. *SIAM review* 51(4):661–703.
- Coleman, J. S. 1990. Foundations of Social Theory.
- Danescu-Niculescu-Mizil, C.; West, R.; Jurafsky, D.; Leskovec, J.; and Potts, C. 2013. No Country for Old Members: User Lifecycle and Linguistic Change in Online Communities. In *WWW*.
- Ducheneaut, N.; Yee, N.; Nickell, E.; and Moore, R. J. 2007. The Life and Death of Online Gaming Communities. In *CHI*.
- Dunbar, R. I. M. 1992. Neocortex Size as a Constraint on Group Size in Primates. *Journal of Human Evolution* 22(6):469–493.
- Edwards, J. 2009. The Genealogy of Communities. <https://www.genealogytoday.com/articles/reader.mv?ID=2885>. Accessed: 2017-10-28.
- Feick, L. F., and Price, L. L. 1987. The Market Maven: A Diffuser of Marketplace Information. *Journal of Marketing* 51(1):83–97.
- Fink, C.; Schmidt, A.; Barash, V.; Cameron, C.; and Macy, M. 2015. Complex Contagions and the Diffusion of Popular Twitter Hashtags in Nigeria. *Social Network Analysis and Mining* 6(1):1–19.
- Fleming, L.; Colfer, L.; Marin, A.; and McPhie, J. 2007. Why the Valley Went First: Aggregation and Emergence in Regional Inventor Networks. In *The Emergence of Organizations and Markets*. 520–544.
- Gaffney, D., and Matias, J. N. 2018. Caveat Emptor, Computational Social Science: Large-Scale Missing Data in a Widely-Published Reddit Corpus. *arXiv preprint arXiv:1803.05046*.
- Girvan, M., and Newman, M. E. J. 2002. Community Structure in Social and Biological Networks. *Proceedings of the National Academy of Sciences* 99(12):7821–7826.
- Goel, S.; Anderson, A.; Hofman, J.; and Watts, D. J. 2015. The Structural Virality of Online Diffusion. *Management Science* 62(1):180–196.
- Hamilton, W. L.; Zhang, J.; Danescu-Niculescu-Mizil, C.; Jurafsky, D.; and Leskovec, J. 2017. Loyalty in Online Communities. In *ICWSM*.
- Hessel, J.; Tan, C.; and Lee, L. 2016. Science, AskScience, and BadScience: On the Coexistence of Highly Related Communities. In *ICWSM*.
- Hirschman, A. O. 1970. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*, volume 25. Harvard university press.
- Kairam, S. R.; Wang, D. J.; and Leskovec, J. 2012. The Life and Death of Online Groups: Predicting Group Growth and Longevity. In *WSDM*.
- Kim, A. J. 2000. *Community Building on the Web: Secret Strategies for Successful Online Communities*. Addison-Wesley Longman Publishing Co., Inc., 1st edition.
- Kossinets, G., and Watts, D. J. 2006. Empirical analysis of an evolving social network. *Science* 311(5757):88–90.
- Kraut, R. E.; Resnick, P.; Kiesler, S.; Burke, M.; Chen, Y.; Kittur, N.; Konstan, J.; Ren, Y.; and Riedl, J. 2012. *Building successful online communities: Evidence-based social design*. Mit Press.
- Kumar, S.; Hamilton, W. L.; Jurafsky, D.; and Leskovec, J. 2018. Community Interaction and Conflict on the Web. In *WWW*.
- Lewin, K. 1951. *Field Theory in Social Science*. Harper.
- Liben-Nowell, D., and Kleinberg, J. 2008. Tracing information flow on a global scale using Internet chain-letter data. *Proceedings of the National Academy of Sciences* 105(12):4633–4638.
- Padgett, J. F., and Powell, W. W., eds. 2012. *The Emergence of Organizations and Markets*. Princeton University Press.
- Pavlopoulou, M. E. G.; Tzortzis, G.; Vogiatzis, D.; and Paliouras, G. 2017. Predicting the evolution of communities in social networks using structural and temporal features. In *Workshop on Semantic and Social Media Adaptation and Personalization*.
- Prentice, D. A.; Miller, D. T.; and Lightdale, J. R. 1994. Asymmetries in Attachments to Groups and to Their Members: Distinguishing between Common-identity and Common-bond Groups. In *Personality and Social Psychology Bulletin*, volume 20. Sage Publications Sage CA: Thousand Oaks, CA. 484–493.
- Ren, Y.; Kraut, R. E.; and Kiesler, S. 2007. Applying Common Identity and Bond Theory to Design of Online Communities. *Organization Studies* 28(3):377–408.
- Rogers, E. M. 2003. *Diffusion of Innovations*. Free Press.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *WWW*.
- Romero, D. M.; Tan, C.; and Ugander, J. 2013. On the Interplay between Social and Topical Structure. In *ICWSM*.
- Ruvio, A., and Shoham, A. 2007. Innovativeness, Exploratory behavior, Market Mavenship, and Opinion Leadership: An Empirical Examination in the Asian Context. *Psychology & Marketing* 24(8):703–722.
- Singer, P.; Flöck, F.; Meinhart, C.; Zeitfogel, E.; and Strohmaier, M. 2014. Evolution of Reddit: From the Front Page of the Internet to a Self-referential Community? In *WWW (Companion)*.
- Stuart, E. A., and Rubin, D. B. 2008. Best Practices in Quasi-experimental Designs: Matching Methods for Causal Inference. In *Best practices in quantitative methods*. 155–176.
- Tajfel, H. 1982. *Social identity and intergroup relations*. Cambridge University Press.
- Tan, C., and Lee, L. 2015. All Who wander: On the Prevalence and Characteristics of Multi-community Engagement. In *WWW*.
- Uzzi, B.; Mukherjee, S.; Stringer, M.; and Jones, B. 2013. Atypical Combinations and Scientific Impact. *Science* 342(6157):468–472.
- Vasilescu, B.; Serebrenik, A.; Devanbu, P.; and Filkov, V. 2014. How Social Q&A Sites Are Changing Knowledge Sharing in Open Source Software Communities. In *CSCW*.
- Wilcoxon, F. 1945. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1(6):80–83.
- Yang, J., and Leskovec, J. 2013. Overlapping Community Detection at Scale: a Nonnegative Matrix Factorization Approach. In *WSDM*.
- Zhang, J.; Hamilton, W. L.; Danescu-Niculescu-Mizil, C.; Jurafsky, D.; and Leskovec, J. 2017. Community Identity and User Engagement in a Multi-Community Landscape. In *ICWSM*.
- Zhu, H.; Chen, J.; Matthews, T.; Pal, A.; Badenes, H.; and Kraut, R. E. 2014. Selecting an Effective Niche: An Ecological View of the Success of Online Communities. In *CHI*.
- Zhu, H.; Kraut, R. E.; and Kittur, A. 2014. The Impact of Membership Overlap on the Survival of Online Communities. In *CHI*.