

Navigating the massive world of reddit: Using backbone networks to map user interests in social media

Randal S. Olson^{1,*}, Zachary P. Neal²

1 Department of Computer Science & Engineering

2 Department of Sociology

Michigan State University, East Lansing, MI 48824, U.S.A.

*** E-mail: olsonran@msu.edu**

Abstract

In the massive online worlds of social media, users frequently rely on organizing themselves around specific topics of interest to find and engage with like-minded people. However, navigating these massive worlds and finding topics of specific interest often proves difficult because the worlds are mostly organized haphazardly, leaving users to find relevant interests by word of mouth or using a basic search feature. Here, we report on a method using the backbone of a network to create a map of the primary topics of interest in any social network. To demonstrate the method, we build an interest map for the social news web site reddit and show how such a map could be used to navigate a social media world. Moreover, we analyze the network properties of the reddit social network and find that it has a scale-free, small-world, and modular community structure, much like other online social networks such as Facebook and Twitter. We suggest that the integration of interest maps into popular social media platforms will assist users in organizing themselves into more specific interest groups, which will help alleviate the overcrowding effect often observed in large online communities.

Introduction

In the past decade, social media platforms have grown from a pastime for teenagers into tools that pervade nearly all modern adults' lives [1]. Social media users typically organize themselves around specific interests, such as a sports team or hobby, which facilitates interactions with other users who share similar interests. For example, Facebook users subscribe to topic-specific "pages" [2], Twitter users classify their tweets using topic-specific "hashtags" [3], and reddit users post and subscribe to topic-specific sub-forums called "subreddits" [4].

These interest-based devices provide structure to the growing worlds of social media, and are essential for the long-term success of social media platforms because they make these big worlds feel small and navigable. However, navigation of social media is challenging because these worlds do not come with maps [5, 6]. Users are often left to discover pages, hashtags, or subreddits of interest haphazardly, by word of mouth, following other users' "votes" or "likes", or by using a basic search feature. Owing to the scale-free structure of most online social networks, these elementary navigation strategies result in users being funnelled into a few large and broad interest groups, while failing to discover more specific groups that may be of greater interest [7, 8].

In this work, we combine techniques for network backbone extraction and community detection to construct a roadmap that can assist social media users in navigating these interest groups by identifying related interest groups and suggesting them to users. We implement this method for the social news web site reddit [9], one of the most visited social media platforms on the web [10], and produce an interactive map of all of the subreddits. An interactive version of the reddit interest map is available online [11].

By viewing subreddits as nodes linked by users with common interests, we find that the reddit social media world has a scale-free, small-world, and modular community structure. The scale-free property is the expected outcome of a preferential attachment process and helps explain the challenges of haphazard navigation. Additionally, the small-world property explains how the big world of reddit can seem small

and navigable to users when it is mapped out. Finally, the modular community structure in which narrow interest-based subreddits (e.g., dubstep or rock music) are organized into broader communities (e.g., music) allows users to easily identify related interests by zooming in on a broader community. We suggest that the integration of such interest maps into popular social media platforms will assist users in organizing themselves into more specific interest groups, which will help alleviate the overcrowding effect often observed in large online communities [4].

Further, this work releases and provides an overview of a data set of over 850,000 anonymized reddit user's interests, thus establishing another standard real-world social network data set for researchers to study. This is useful because, although reddit is among the largest online social networks and has been identified as a starting point for the viral spread of memes and other online information [12], it has been relatively understudied [4, 13, 14]. This data set can be downloaded online at [15].




Figure 1. Reddit interest network. The largest components of the reddit interest network is shown with 10 interest meta-communities annotated; it closely matches the structure of other online social networks including Flickr and Yahoo360 [16]. Each node is a single subreddit, where color indicates the interest meta-community that the subreddit is a member of. Nodes are sized by their weighted PageRank to provide an indication of how likely a node is to be visited, and positioned according to the OpenOrd layout in Gephi to place related nodes together. An interactive version of the reddit interest map is available online at <http://rhiever.github.io/redditviz/clustered/>

Results

Reddit interest map

For the final version of the reddit interest map, we use the backbone network produced with $\alpha = 0.05$ (see Methods). This results in a network with 59 distinct clusters, which we call *interest meta-communities*. In Figure 1, the nodes (i.e., subreddits) are sized by their weighted PageRank [17] to provide an indication of how likely a node is to be visited, and positioned according to the OpenOrd layout in Gephi [18] to place related nodes together.

Through this method, we immediately see several distinct interest meta-communities, 10 of which are annotated in Figure 1. These interest meta-communities act as starting points in the interest map to show the broad interest categories that the entire reddit community is discussing. From these starting points, users can zoom in on a single broad interest category to find subreddits dedicated to more specific interests, as shown in Figure 2. Notably, there is a large, orange interest meta-community in the center of the interest map that overlaps with several other interest meta-communities. This orange interest meta-community represents the most popular, general interest subreddits (e.g., “pictures” and “videos”) in which users of all backgrounds regularly participate, and thus are expected to have considerable overlap with many other communities.

Figure 2 depicts zoomed-in views of two interest meta-communities annotated in Figure 1. In Figure 2A, the “sports” meta-community, specific sports teams are organized around the corresponding sport that the teams play in. For example, subreddits dedicated to discussion of the Washington Redskins or Denver Broncos – relatively small, specific subreddits – are organized around the larger, more general interest NFL subreddit where users discuss the latest NFL news and games. Similarly in Figure 2B, the “programming” meta-community, subreddits dedicated to discussing programming languages such as Python and Java are organized around a more general programming subreddit, where users discuss more general programming topics.

This backbone network structure naturally lends itself to an intuitive interest recommendation system. Instead of requiring a user to provide prior information about their interests, the interest map provides a hierarchical view of all user interests in the social network. Further, instead of only suggesting interests immediately related to the user’s current interest(s), the interest map recommends interests that are potentially two or more links away. For example in Figure 2A, although the Miami Heat and Miami Dolphins subreddits are not linked, Miami Heat fans may also be fans of the Miami Dolphins. A traditional recommendation system would only recommend NBA to a Miami Heat fan, whereas the interest map also recommends the Miami Dolphins subreddit because they are members of the same interest meta-community.

Network properties

In Figure 3, we show a series of network statistics to provide an overview of the backbone reddit interest network. These network statistics are plotted over a range of α cutoff values for the backbone reddit interest network (see Methods) to demonstrate that the interest network we chose in Figure 1 is robust to relevant α cutoff values.

As expected, the majority of the edges are pruned by an α cutoff of 0.05 (Figure 3, top left). This result demonstrates that the backbone interest network is stable with an α cutoff ≤ 0.05 , which is the most relevant range of α cutoffs to explore. Surprisingly, 80% of the subreddits that we investigated – roughly 12,000 subreddits – do not have enough users that consistently post in another subreddit to maintain even a single edge with another subreddit. The majority of these 12,000 subreddits likely do not have any significant edges due to user inactivity, e.g., some subreddits have only a single user that frequently posts to them (Table 1). Another factor that likely contributes to the 12,000 unlinked subreddits is temporary interests, i.e., an interest such as the U.S. Presidential election that temporarily

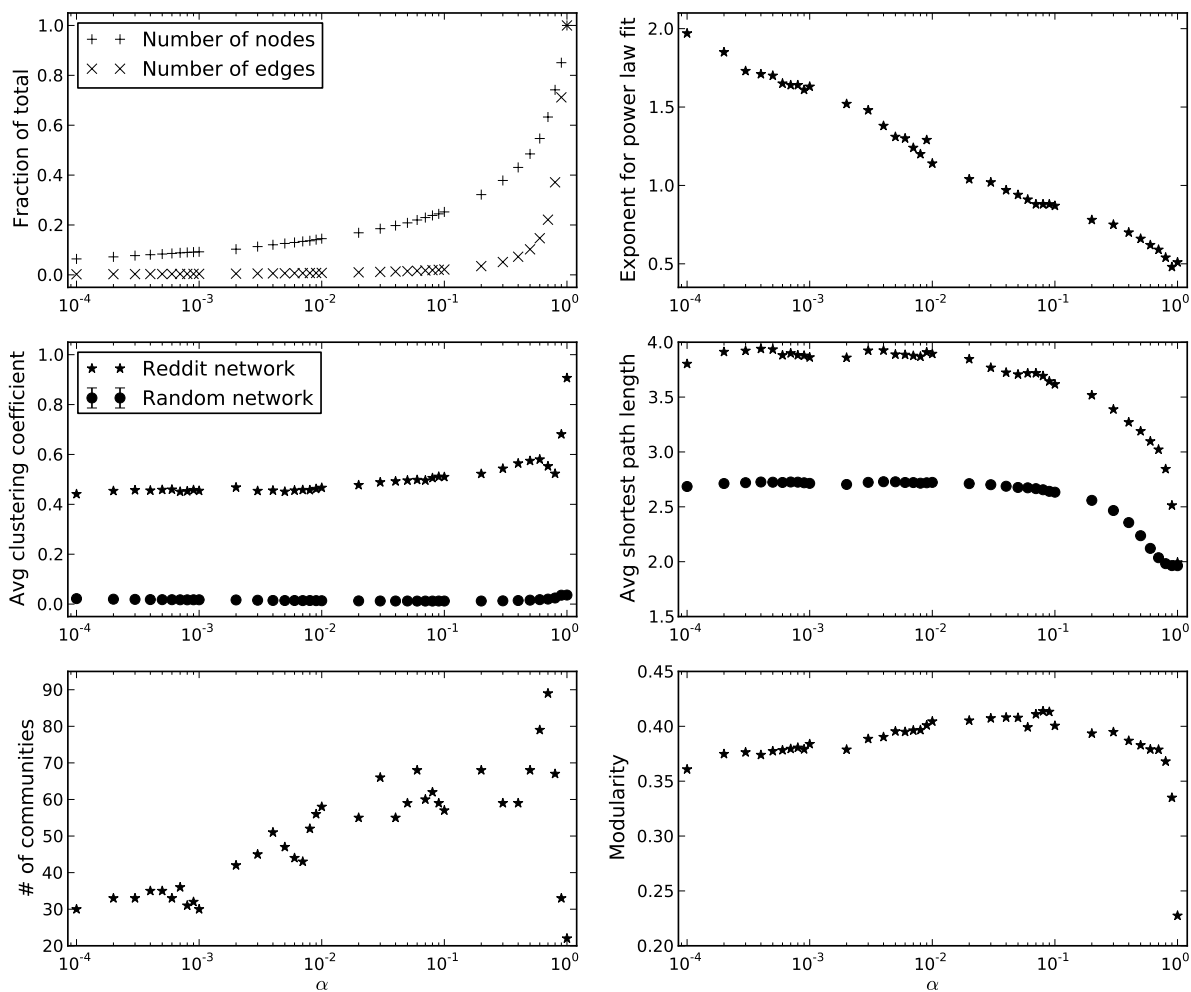


Figure 3. Network statistics for the backbone network. Sensitivity analysis of the reddit interest network over a range of α cutoff values. Lower α means that fewer statistically significant edges are pruned. In general, this sensitivity analysis shows that the backbone interest network is stable for α cutoff values ≤ 0.05 . Error bars for the Erdős-Rényi random networks are two standard deviations over 30 random networks, and are too small to show up on the graph. Note the logarithmic scale of the x-axis.

draws a large number of people together, but eventually fades into obscurity again.

Next, we are interested in exploring whether the backbone reddit interest network is a scale-free network, where preferential attachment to subreddits results in a few extremely popular (i.e., connected) subreddits and mostly unpopular subreddits. As such, scale-free networks are known to have node degree distributions that fit a power law [7, 8]. Regardless of the α cutoff, we observed that the node degree distribution of all backbone reddit interest networks fit a power law ($R^2 \approx 0.91$ for $k \geq 50$; Figure 3, top right). This scale-free network structure is likely partially due to reddit’s default subreddit system [19], where newly registered users are subscribed to a set of 20 subreddits by default.

Furthermore, we want to confirm that the backbone reddit interest network is a small-world network [20]. Small-world networks are known to contain numerous clusters, as indicated by a high average clustering coefficient, with sparse edges between those clusters, which results in an average shortest path length between all nodes (L_{sw}) that scales logarithmically with the number of nodes (N):

$$L_{sw} \approx \log_{10}(N) \tag{1}$$

Figure 3 (middle left and middle right) depicts the average clustering coefficient and shortest path length for all nodes in the backbone reddit interest network. Compared to Erdős-Rényi random networks with the same number of nodes and edges, the backbone network has a significantly higher average clustering coefficient. Similarly, the measured average shortest path length of the backbone network (α cutoff = 0.05) follows Equation 1, with $L_{sw} = \log_{10}(2,347) = 3.37 \approx 3.71$ from Figure 3 (middle right). Thus, the backbone reddit interest backbone network qualitatively appears to exhibit small-world network properties.

To quantitatively determine whether the reddit interest network exhibits small-world network properties, we used the small-worldness score (S_G) proposed in [21]:

$$S_G = \frac{C_G/C_{rand}}{L_G/L_{rand}} \tag{2}$$

where C is the average clustering coefficient, L is the average shortest path length between all nodes, G is the network the small-worldness score is being computed for, and “rand” is an Erdős-Rényi random network with the same number of nodes and edges as G . If $S_G > 1$, then the network is classified as a small-world network. For the backbone reddit network, we calculated $S_G = 14.2$ ($P < 0.001$), which indicates that the reddit interest network exhibits small-world network properties.

Now that we know that the backbone reddit interest network is scale-free and exhibits small-world network properties, we want to study the community structure of the backbone network. Shown in Figure 3 (bottom right), the backbone network exhibits a consistently high modularity score with an α cutoff as high as 0.9, implying that even a slight reduction in the number of edges in the backbone network reveals the reddit interest community structure. Correspondingly, depicted in Figure 3 (bottom left), the number of identified communities (i.e., clusters) remains relatively low until the α cutoff is reduced to ≤ 0.9 . As the α cutoff is reduced, the number of identified communities generally decreases, which coincides with the loss of nodes as α decreases. Thus, the backbone reddit interest network has ≈ 30 core communities, and another ≈ 30 weakly linked communities that are lost as a more stringent α cutoff is applied.

Discussion

We have shown that backbone networks can be used to map and navigate massive interest networks in social media. By viewing the big world of reddit as a hierarchical map, users can now explore related interests without providing any prior information about their own interests. Future applications of this method may also facilitate navigation of other popular social network platforms such as Facebook and Twitter.

Furthermore, such an interest map could allow social media users to self-organize into more specific interest forums, thus reducing preferential attachment to large, general interest forums and alleviating the issues that arise in overcrowded social network forums [4]. Given previous work that suggests network properties such as small-worldness and even modularity can result solely from network growth processes [22], it would be interesting in future work to observe what processes govern network growth when users have access to an interests map like those shown in Figures 1 and 2, and what network properties emerge from these growth processes.

This work provides a unique view of reddit that debunks a common misconception of the social news web site. Typically, outsiders view reddit as a single, homogeneous entity that acts as one, e.g. “Should Reddit Be Blamed for the Spreading of a Smear?” [23]. In contrast, the reddit interest map shown here provides a different view of reddit, where many users organize themselves into cliques based on shared interests and rarely interact with other reddit users outside their clique. In that light, we hope this work reveals that, like many social communities (online or offline), reddit is a community composed of a diverse group of people that are brought together by thousands of seemingly-unrelated interests.

Additionally, we explored the network properties of the backbone reddit interest network that we composed from the posting behavior of over 850,000 active reddit users. In this analysis, we found that the reddit interest network has a scale-free, small-world, and modular community structure, corroborating findings in many other online social networks [24, 25]. Uniquely, reddit potentially enforces a scale-free network structure on its users by automatically subscribing all new users to the same set of 20 subreddits [19]. Exploring the effect of automatically subscribing users to a fixed set of interest-specific forums on social interest network structure could be another interesting venue of future work. To expedite future analyses of the reddit interest network, we have provided the raw, anonymized data set available to download online [15].

It is important to note that the sample of user behavior we have taken is cross-sectional, reflecting users’ reddit posts and thus the relationships among reddit interests at a fixed point in time in mid-2013. However, as users’ interests evolve, so too do the relationships among them [26]. In some cases, highly specialized and related subreddits may fuse into a single subreddit, while in other cases a general subreddit may split into multiple more specialized ones. Thus, such an interest map would require periodic (or, ideally, real-time) updating to accurately reflect dominant interests in the social network and their relationships to one another.

Methods

To acquire the data for this study, we mined user posting behavior data from reddit by first gathering the user names of 876,961 active users that post to 15,122 distinct subreddits (see Table S1 for more detail). We note that reddit reports to have over 2.6 million registered users as of December 2013 [27], so this data set represents a random sample of roughly 1/3 of the total active users on reddit. For each of the users, we gathered their 1,000 most recent link submissions and comments, counted how many times they post to each subreddit, and registered them as interested in a subreddit only if they posted there at least 10 times. We applied this threshold of at least 10 posts to filter out users that are not active in a particular subreddit.

From these data, we defined a bipartite network \mathbf{X} , where $X_{ij} = 1$ if user i is an active poster in subreddit j and otherwise is 0. We then projected this as a weighted unipartite network \mathbf{Y} as $\mathbf{X}\mathbf{X}'$, where Y_{ij} is the number of users that post in both subreddits i and j . This resulted in 4,520,054 non-zero edges between the subreddits. Details of the raw weighted subreddit network are shown in Table 1.

Due to the challenges associated with analyzing large weighted networks, we reduced the number of edges in the weighted subreddit network using a backbone extraction algorithm [28]. This backbone extraction algorithm preserves edges whose weight is statistically incompatible, at a given level of significance α , with a null model in which edge weights are distributed uniformly at random. In the resulting

Table 1. Edge weights in the raw and backbone reddit interest networks

Network	Mean	Minimum	Maximum
Raw	17.77	1	309,985
Backbone	0.0052	0.00068	0.1997

backbone network, two subreddits are linked if the number of users who post in both of them is statistically significantly larger than expected in a null model, from the perspective of *both* subreddits. To combine the directed edges between each two nodes, we replaced the two directed edges with a single undirected edge whose weight is the average of the two directed edges.

Thus, this technique defines a network of subreddit pathways along which there is a high probability users might traverse if they navigate reddit by following the posts of other users. Adjusting the α parameter allows the backbone network to include more (e.g., when α is larger) or fewer (e.g., when α is smaller) such pathways. Figure 3 summarizes the topological properties of backbones extracted using a range of α parameter values; in the findings and discussion we focus on a backbone extracted using the conventional $\alpha = 0.05$.

We used Python’s PRAW package¹ to gather the data and Python’s NetworkX package [29] to compute all network statistics. In the backbone graph, we focus only on the largest connected component. We detected network communities using [30] and visualized the communities using the OpenOrd node layout, both as implemented in Gephi [18].

Acknowledgments

We gratefully acknowledge the support of the Michigan State University High Performance Computing Center and the Institute for Cyber Enabled Research (iCER). We thank Arend Hintze, Christoph Adami, and Emily Weigel for helpful feedback during the preparation of this manuscript.

References

1. Rainie L, Wellman B (2012) Networked: The new social operating system. The MIT Press.
2. Strand JL (2011) Facebook: Trademarks, fan pages, and community pages. Intellectual Property & Technology Law Journal 23: 10–13.
3. Chang HC (2010) A new perspective on twitter hashtag use: diffusion of innovation theory. Proceedings of the American Society for Information Science and Technology 47: 1–4.
4. Gilbert E (2013) Widespread underprovision on reddit. In: Proceedings of the 2013 conference on Computer supported cooperative work. New York, NY, USA: ACM, CSCW ’13, pp. 803–808. doi:10.1145/2441776.2441866.
5. Boguna M, Krioukov D, Claffy KC (2008) Navigability of complex networks. Nature Physics 5: 74–80.
6. Benevenuto F, Rodrigues T, Cha M, Almeida V (2012) Characterizing user navigation and interactions in online social networks. Information Sciences 195: 1–24.
7. Albert R, Jeong H, Barabási AL (1999) Internet: Diameter of the world-wide web. Nature 401: 130–131.

¹Python Reddit API Wrapper (PRAW): <https://github.com/praw-dev/praw>

8. Barabási AL, Albert R, Jeong H (2000) Scale-free characteristics of random networks: the topology of the world-wide web. *Physica A: Statistical Mechanics and its Applications* 281: 69–77.
9. reddit (2013). What is reddit? URL http://www.reddit.com/wiki/faq#wiki_what_is_reddit.3F.
10. Alexa (2013). reddit alexa ranking. URL <http://www.alexa.com/siteinfo/reddit.com>.
11. Olson RS (2013). redditviz, the interactive reddit interest map. URL <http://rhiever.github.io/redditviz/clustered/>.
12. Sanderson B, Rigby M (2013) We've reddit, have you? What librarians can learn from a site full of memes. *College & Research Libraries News* 74: 518–521.
13. Wasike BS (2011) Framing social news sites: An analysis of the top ranked stories on reddit and Digg. *Southwestern Mass Communication Journal* 27.
14. Merritt E (2012) An Analysis of the Discourse of Internet Trolling: A Case Study of Reddit.com. Ph.D. thesis.
15. Olson RS (2013). reddit user posting behavior (mid-2013). URL <http://dx.doi.org/10.6084/m9.figshare.874101>.
16. Kumar R, Novak J, Tomkins A (2010) Structure and evolution of online social networks. In: *Link Mining: Models, Algorithms, and Applications*, Springer. pp. 337–357.
17. Page L, Brin S, Motwani R, Winograd T (1999) The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab.
18. Bastian M, Heymann S, Jacomy M (2009) Gephi: An open source software for exploring and manipulating networks. In: Adar E, Hurst M, Finin T, Glance NS, Nicolov N, et al., editors, ICWSM. The AAAI Press.
19. reddit (2011). Saying goodbye to an old friend and revising the default subreddits. URL <http://blog.reddit.com/2011/10/saying-goodbye-to-old-friend-and.html>.
20. Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286: 509-512.
21. Humphries MD, Gurney K (2008) Network small-world-ness: A quantitative method for determining canonical network equivalence. *PLoS ONE* 3: e0002051.
22. Hintze A, Adami C (2010) Modularity and anti-modularity in networks with arbitrary degree distribution. *Biology Direct* 5: 32+.
23. Kang JC (2013). The New York Times: Should reddit be blamed for the spreading of a smear? URL <http://www.nytimes.com/2013/07/28/magazine/should-reddit-be-blamed-for-the-spreading-of-a-smear.html?pagewanted=all>.
24. Ahn YY, Han S, Kwak H, Moon S, Jeong H (2007) Analysis of topological characteristics of huge online social networking services. In: *Proceedings of the 16th International Conference on World Wide Web*. New York, NY, USA: ACM, WWW '07, pp. 835–844. doi:10.1145/1242572.1242685.
25. Mislove A, Marcon M, Gummadi KP, Druschel P, Bhattacharjee B (2007) Measurement and analysis of online social networks. In: *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*. New York, NY, USA: ACM, IMC '07, pp. 29–42. doi:10.1145/1298306.1298311.

26. Banerjee N, Chakraborty D, Dasgupta K, Mittal S, Joshi A, et al. (2009) User interests in social media sites: An exploration with micro-blogs. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management. New York, NY, USA: ACM, CIKM '09, pp. 1823–1826. doi:10.1145/1645953.1646240.
27. reddit (2013). About reddit. URL <http://www.reddit.com/about>.
28. Serrano M, Bogu M, Vespignani A (2009) Extracting the multiscale backbone of complex weighted networks. Proceedings of the National Academy of Sciences 106: 6483-6488.
29. Hagberg AA, Schult DA, Swart PJ (2008) Exploring network structure, dynamics, and function using NetworkX. In: Proceedings of the 7th Python in Science Conference (SciPy2008). Pasadena, CA USA, pp. 11–15.
30. Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. Journal of Statistical Mechanics: Theory and Experiment 2008: P10008.

Supplementary Information

Table S1. Descriptive statistics of the bipartite (user-to-subreddit) network

Statistic	Value
Total # of users	876,961
Total # of subreddits	15,122
Average # of subreddits per user	9.69
Minimum # of subreddits per user	1
Maximum # of subreddits per user	112
Average # of users per subreddit	561.8
Minimum # of users per subreddit	1
Maximum # of users per subreddit	523,025