

# Analyzing Social Support on the Experience Project

## CS 224W Project

Desmond C. Ong (dco@stanford) & Shen Minghan (sminghan@stanford) (Group 1)

### Abstract

More and more people are turning online for social support, and large social networking/online support sites such as the Experience Project have seen incredible growth in the past few years. Using network analysis, we analyze the structure and dynamics of social support in the Experience Project, where users anonymously write posts about emotional events and comment on each others posts. We define two types of social support: *active* suggestions or solutions to their problem, and *passive* sympathy and commiseration, and analyze the differences in propagation behavior in the network, finding some support for the hypothesis that active support results in longer cascades. We find that social support improves network modularity, and that influence on cascades is not strongly correlated with number of friends.

### Introduction

With the increasing ubiquity of social networks, more and more people are turning to the internet for social interactions, and increasingly, to seek **social support**, especially when they are unable to find it in real life. Participation in online forums for specific (mental and physical) illnesses [1], for example, can result in positive mental outcomes [2] such as empowerment against the illness [3]. The Experience Project (EP) is a social networking site that started out as a support group for multiple sclerosis patients, but has since expanded to include support groups for any type of shared experiences. It has since grown exponentially in the past decade, boasting over 34 million experiences shared by users (as of Nov 2013) since its inception in late 2006. Users on EP share very personal stories, sometimes just writing to share, which is known to have therapeutic outcomes [4], and other times actively seeking comfort and advice to solve a problem. The community nature of the website encourages and facilitates interactions between users that are emotionally rich and constructive. As more and more people are living their lives online on websites such as EP, baring their souls and seeking friendship and support on the Web, it is important to study the nature and dynamics, and the promises and potential perils, of online social support.

There have been numerous studies on the benefits of social support both for the support seeker and provider (see [5, 6, 7, 8] for extensive reviews); here, we focus instead on the dynamics of online social support. How can we measure online social support, and how does that support propagate in an online social network? In particular, are there different types of social support that are often given in online scenarios, and do these types of social support differ in how they spread?

In this work, we propose using simple NLP tools to characterize two different types of social support on EP: what we define as “**active support**” and “**passive support**”. Active

support is concrete advice or suggestions given by the user meant to actively help the support seeker deal with the problem, while passive support includes sympathetic notes of understanding. We seek to concretely test the hypothesis that these two different types of social support have different dynamics: active support would encourage a “pay it forward” behavior, resulting in large cascades, while passive support would encourage commiseration, or encourage the formation of clusters and communities in the network.

### Review of prior work

Previous work in analyzing the dynamics of sentiment or emotion have looked at the propagation through online media such as blogs [9, 10] and Twitter [11]. [9] used sentiment analysis to analyze sentiment cascade flow in hyperlink networks, where nodes were blogs and an edge from A to B represents that blog A cited blog B. They used a lexicon-based bag-of-words model by building off the Harvard Inquirer and SentiWordNet, two large sentiment-annotated lexica, added modifications such as adding emoticons and negation, to label the sentiment of their blog posts. They characterized shallow cascades as having only small changes in subjectivity from the parent post, while deeper cascades have large polarizing changes in subjectivity and sentiment (possibly due to controversial commenting/linking on different blogs).

[10] took a similar approach by analyzing sentiment propagation on LiveJournal, a blogging website where bloggers (nodes) can be friends (edge) with other bloggers on the site. They characterized the mood of posts, and looked for similarities in moods with other blogs with whom the original blogger is friends with. They defined a temporally graded sentiment propagation, where sentiment has propagated if the overall mood of posts of a user at some time  $t_j > t_i$  is closer to the mood of the original post made at time  $t_i$ , and this influence decays as  $t_j$  increases. They find that users who have a cascading effect of sentiment propagation tend to have higher degree (i.e. more friends), have friends who themselves have high degree, and make fewer (but more high-impact) posts.

A third study [11] studied Subjective Well-Being (SWB) in a large Twitter network using OpinionFinder, another lexicon-based sentiment analysis software. They raise an interesting challenge for research in this field in trying to disambiguate between assortive mixing or homophilic attachment (do happy people seek out other happy people?), and contagion (do happy people cause other people to be happy). They report pair-wise and neighborhood level SWB assortativity, and find significant levels of assortativity at both levels, i.e. friends and neighbors of a Twitter user tend to have the same level of SWB, which may be a close proxy to

mood/emotion/sentiment assortativity.

Thus, while our paper is looking at an unrelated phenomenon (social support) within a social networking site, we draw inspiration from previous work in using lexicon-based sentiment analysis [9, 11], consider dynamics in temporally-graded propagation [9, 10], and looking at assortativity [11].

## Data

The corpus used in this paper is from the Experience Project (EP), which is a social networking site where users join groups based on shared experiences, e.g. “I am a military wife”, or “I am a cancer survivor”, or even “I am lonely”. Within these groups, they can post stories, and other users can comment on them. This provides a large corpus of richly emotional text, and, moreover, chronicles social interactions between users. Previous sentiment analysis work using the same dataset has been done by [12] and [13]. The original dataset from [12] had 22k stories, and we collected an additional 30k stories; the stories were initially posted on the website from Feb, 2006 through June, 2010.

The data that we used consists of users who write posts (which we call the original poster, OP), and other users who comment on those posts, along with the timestamp and text content of the post. We defined nodes as users on the network, and we define an edge from a user to an OP if a user wrote a comment on an OP’s story – an edge from A to B indicates “A commented on B’s story”. Thus, an edge represents an interaction, and we can study the nature and dynamics of these interactions.

## Summary Statistics of (Static) Dataset

The static network that we constructed based on the data had the following structural properties (in Table 1). By static, we mean that we removed the time information from the edges, i.e. if A had comment on B 10 times, we only count a single edge now between A and B, which reduced the total number of observations we had (from 153,707 time-accounted edges). The static graph has 36k nodes and 80k edges between them. Due to the way that we constructed the network, there are no zero degree nodes. A node with zero in degree indicates a user who only comments on other users’ posts, but received no comments on their own posts (and these are not added to our network), while a node with zero out degree indicates a user who has only received comments on their posts, but did not comment on others. There is also a large amount of bidirectional interactions, where A comments on B’s post and B comments on A’s post, with almost an eighth of the total edges being bidirectional edges. The largest weakly connected component was of size 33k (90.6% of the graph), which shows that the large majority of the community are interacting with each other: the strongly connected component size is about 3.5k, or 9.5% of the graph. The average clustering coefficient, after [14], is:

$$C = \frac{1}{N} \sum_i C_i; \quad C_i = 2 \frac{e_i}{k_i(k_i - 1)} \quad (1)$$

where  $N$  is the total number of nodes,  $C_i$  is the clustering coefficient of node  $i$  with degree  $k_i$ ,  $e_i$  is the number of edges between the neighbors of node  $i$ , and the average clustering coefficient is simply the average of all the clustering coefficients in the network. We can also calculate the modularity coefficient, after [15], given as:

$$Q = \frac{1}{2m} \sum_{s \in S} \sum_{\{i,j\} \in s} \left[ A_{ij} - \frac{k_i k_j}{2m} \right] \quad (2)$$

where  $A_{ij}$  equals the weight of edge  $(i, j)$ , and  $k_i = \sum_j A_{ij}$  is the sum of the weights on the outgoing edges from  $i$ , and  $m = \frac{1}{2} \sum_{i,j} A_{ij}$  is half the total sum of weights in the network. The sum is done over all  $(i, j)$  pairs in the same community  $s$ .

Nodes	36,712
Edges	80,148
Zero Deg Nodes	0
Zero InDeg Nodes	23,876
Zero OutDeg Nodes	5,236
NonZero In-Out Deg Nodes	7,600
Unique directed edges	80,148
Unique undirected edges	77,986
BiDir Edges	10,122
Weakly connected component size	33,255 (90.6%)
Strong connected component size	3,502 (9.5%)
Clustering Coefficient	0.0270
Modularity $Q$	0.605
Approx. full diameter	14
90% effective diameter	6.122187

Table 1: Basic structural properties of the directed static EP network

## Model and Algorithm

### Defining types of social support

There is an extensive psychological literature on social support (e.g. [5]), mostly focused on social support between friends and family in real life. A common distinction is into four types of social support: **Emotional** (expressing care/love), **Instrumental** (offering some kind of tangible support, such as financial support, or helping someone with work), **Informational** (offering advice), and **Appraisal/Esteem** (bolstering the person’s self-confidence, or helping to reframe the situation, such as “looking on the bright side”). We chose not to go with these distinctions on EP because they are not appropriate to characterize these online interactions. For one, very few people give instrumental support on online websites, and for another, interactions are not face-to-face and often with strangers, which limits the amount of empathy and emotional support people provide to each other.

A third, more important reason, is that there is a large proportion of interaction on EP where users share their own experiences (hence the name of the site), which results in a self-focused reply or interaction, rather than other-focused on the

original poster (OP). For example, in some responses to a post describing a breakup (relationships are the most common topic on EP), other users might offer comments such as “*I feel your pain! I’ve been through it before*”; these posts mainly consist of recounting their own stories, in a sense commiserating with the OP. This experience sharing is a different type of support that does not fall into the traditional academically defined categories of social support, but is equally important and drives both offline and online social support groups.

For the scope of this work, we defined two types of social support.

- The first, is **active support**, whereby users offer advice, and reframing of the situation (a combination of the Informational and Appraisal/Esteem categories), i.e. offering active solutions to help the OP develop solutions to their problems.
- The other, **passive support**, is characterized by users describing their own experiences, or offering sympathy, and understanding (which includes some aspects of Emotional Support).

In order to define these constructs within the data, we consulted validated lexicons, such as AFINN [16], the Affective Norms for English Words (ANEW) [17], and the Linguistic Inquiry and Word Count (LIWC) [18] to create a set of uni-, bi-, and trigram lexicons to classify active and passive comments. In addition, we incorporated insights from the psychological literature; for example, increased use of first-person pronouns in discourse is related to more self-focused thinking, while increased use of second and third-person pronouns is associated with other-focused thinking [19], and hence we added first-person pronouns to the passive support lexicon, and second-person pronouns to the active support lexicon, although we down weighted these pronouns relative to other words in the lexicon due to the relative frequency of these pronouns. Both of the final lexicons we used have a total of 75 words each.

### Cascade Algorithm

Our cascade algorithm uses Depth-First Search (DFS) that uses temporal order and the classified support levels to enforce edge criteria.

For temporal order, we use the timestamps for each comment and add it to the edge data, and in the DFS only allow edge traversal when the timestamp of the candidate edge is after the immediately preceding edge. Learning from [10], we restricted the time relevancy of each comment to seven days, as after the period of a week it is unlikely that a previous comment would be influential on the current one.

An additional edge criterion is used to compare the effects of having comments that are more actively supportive than passively supportive, and vice versa. Each comment is passed into our classifier and given an ActiveScore and PassiveScore, and the results are also appended to the edges, which we use

as another criteria for edge traversal, described in the next section.

We allowed for multiple edges with unique timestamps between pairs of nodes, which is common when users engage in conversation in the comments section or over multiple stories (nodes). We did not want conversations to show up as long cascades, since that is not the propagation of support we are trying to detect, so a final criteria was added that only allows travelling to a node that is not yet in the current cascade. This is achieved by passing the visited set to each level of the DFS.

In our initial trial runs we found that a slightly larger proportion of the edges pass the criteria for the Active subgraph, which leads to a larger number of candidate starting nodes for cascade generation. To get a fair comparison we ran the cascade detection on 5000 nodes chosen at random from the set of candidates and computed the total number of cascades and cascade lengths, and took the average of 100 simulations.

## Results

### Assigning sentiment and social support attributes to edges

Recall that an edge from A to B represents a comment that user A made on OP B’s post. We used the lexicon based approach described earlier to assign active and passive scores to each edge. We classified an edge as conveying active support if its active score was greater than its passive score, and its active score was greater than 2 (i.e. it contained more than 2 words in the active lexicon), and similarly, we classified an edge as conveying passive support if its passive score was greater than its active score, and its passive score was greater than 2. We then induced two subgraphs, an Active network and a Passive network, on these two sets of edges, shown in Table 2. Accounting for time, there were a total of 153,707 edges in the total network, 15,419 edges in the active network, and 12,358 edges in the passive network.

	Full	Active	Passive
Nodes	36,712	12,518	11,019
Time-accounted Edges	153,707	15,419	12,358
Static Edges	80,148	13,928	11,250
Clustering Coefficient	0.0270	0.0101	0.0081
Modularity	0.605	0.873	0.884

Table 2: The Full network, and the Active and Passive subgraphs induced by the conditions described earlier.

### Hypothesis 1: Larger clusters and community for Passive support.

We calculated the clustering coefficient and modularity in Table 2, as described earlier. We find that the average clustering coefficient drops for the Active and Passive subgraphs as compared to the full graph, but the modularity coefficient increases for both. Thus, there does not seem to be a difference between the Active and Passive subgraphs, but by imposing the criterion that ActiveScore and PassiveScore must be

greater than 2, eliminates a majority of non-supportive comments (recall that the Full network contains all the comments observed, without filtering for supportive content), we can conclude that both types of social support seem to be associated with improved community structure (increased modularity). This comes at a cost of clustering, which could suggest the formation of smaller, but more well-defined communities. This improved community structure with the social support criterion could be because of homophily, i.e. that people with similar interests join similar communities and are much more likely to give social support to each other, or because of a causal mechanism, i.e. people that receive social support are more likely to join the communities of their benefactors, but it is very hard to differentiate between the two [11].

### Hypothesis 2: Longer cascades for Active support

In order to test the hypothesis that active support causes longer cascades than passive support, because it might encourage people to pay it forward, we implemented our cascade algorithm on the two subgraphs, Active and Passive, with the additional Temporal Order condition that in order for a subsequent edge to be added to a cascade, it has to occur at a later time than the preceding edge. Because the number of nodes in both graphs are slightly unbalanced, we picked a random subset of 5,000 nodes on both graphs, and calculated the number of cascades and cascade lengths for this random subset. We repeated this simulation 100 times. The cascade length distribution is plotted on a semi log plot in Fig. 1, with the mean and standard errors from the values in our simulations. We can see that the cascade distribution in both the Active and Passive subgraphs seem to be indistinguishable at low cascade lengths ( $< 8$ ), but at high cascade lengths, we can see that the Active subgraph decays more slowly. Thus, these results give encouraging results for the hypothesis that Active support might be associated with longer cascades.

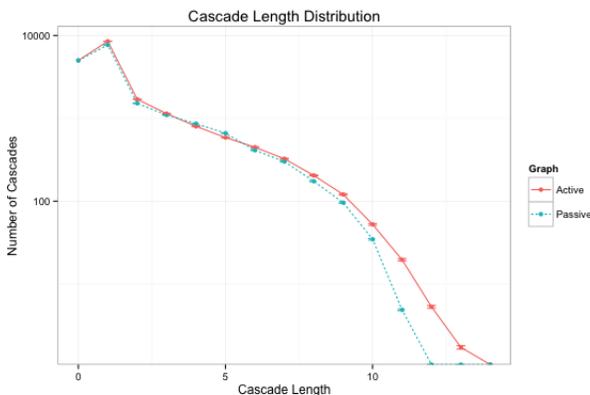


Figure 1: **Cascade Length Distribution, semi log scale.** Cascade lengths for the Active and Passive subgraphs, from 100 simulations of both graphs starting with a random subset of 5,000 nodes. Error bars represent  $\pm 1$  standard error.

### Hypothesis 3: Investigating sources of cascades

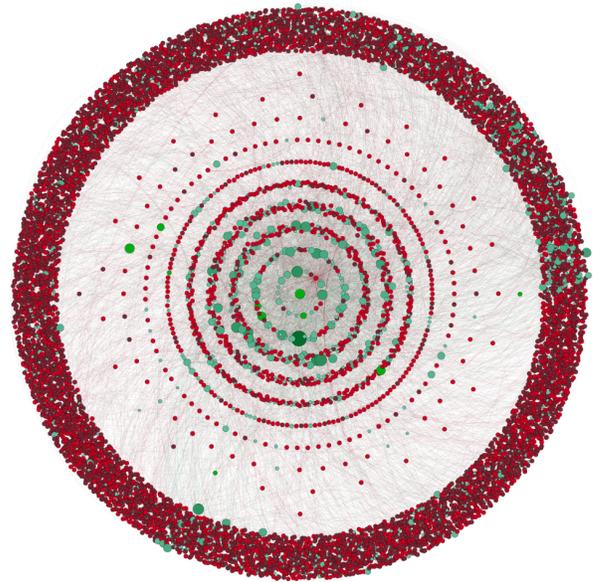


Figure 2: **Top Passive Node.** Node in passive subgraph with the most number of cascades of length 4 or greater is in the center of this figure. Each successive ring is an additional hop away from the start, and the outermost ring is unreachable. Red nodes have out degree zero and the cascade will terminate at red nodes. Figure was created using [20]

An additional hypothesis we had was that sources of cascades would be influential in other ways, such as having more friends, but that this is not automatically guaranteed either, as ultra-social users might not be good supporters.

First, we identified the 15 nodes with the largest number of cascades (greater than 4; we wanted to discount short cascades), in both graphs. By searching of the average number of friends of these 30 nodes is 89, and the top passive node (with 143 cascades) has 152 friends, while the top active node (with 201 cascades) only has 20 friends. Thus, while it seems that on average, the influential nodes (sources of cascades) have a higher number of friends, this might not be true for everyone. One problem is the definition of friendship on this site: A can choose to add B in his circle, in which case A is a “fan” of B and B is a “favorite” of A. If B reciprocates by adding A to his circle as well, then they both become “friends”. Unfortunately, we cannot determine from the website how many fans a particular user has (it is private information), so it might very well be the case that these top sources of cascades might attract many fans, but might not convert those fans into friends.

In addition, the influence of a node should be related to, but not entirely determined by out degree. Both the top passive and the top active node have an out degree of 18, which is not surprising given that a source of cascades needs to have a high out-degree to influence others. However, as a counter-

example, the node with the highest out-degree in the passive subgraph has out degree 37 (and 41 in the active subgraph), but had less than 10 cascades in both subgraphs. From checking the site, this user has joined over 65k groups and has over 22k friends! Thus, our results seem to suggest that while the ability to cause cascades (and hence influence via giving social support) is definitely dependent on out-degree (if a user does not comment on other users, s/he cannot offer social support), after a certain out degree they become unrelated.

To explore the cascade structure further, we took the top passive node, or the node in the passive subgraph with the highest number of cascades, and plotted the network surrounding him in Fig. 2. The top node is in the center, and it is surrounded by concentric rings, where each successive ring contains nodes that can be reached at increasing number of hops away. The large, outermost ring indicates the nodes that are unreachable. Green nodes indicates nodes with positive out-degree, which allow the cascade to propagate, while red nodes are terminal nodes with no out-degree, and the cascades terminate at them. We can visualize the propagation of cascades by noticing that as one gets further away from the source node, the proportion of red nodes increase. We additionally visualize the top active node in Fig. 3, and we indicate an example of a long cascade.

### Conclusion

In this work, we characterized two different types of social support: active support which consists of concrete advice or suggestions, and passive support which includes sympathy and emotional support. We find, after inducing subgraphs on edges which were offering mainly active support and mainly passive support, that these subgraphs had higher modularity than the full network, which suggests that giving and receiving social support is associated with greater community structure. We find also that Active support tends to result in longer cascades as compared to Passive support, which supports our hypothesis that Active support might result in more “paying it forward” behavior. Finally we find that being a source of a cascade is not highly correlated with number of friends.

There are several limitations to our work. Firstly, the data that we collected is a small subset of all the interactions on the site. Secondly, we did not consider the content of the original post, which might help us to filter by whether the original post is support seeking.

These results have interesting applications to understanding how people interact and give and receive social support in online social communities. Identifying the characteristics of people who are good sources of social support will help the moderators of the website to identify volunteers and encourage them to be more active on the site (a program that EP currently has), and might help to train people to be good sources of social support.

### Acknowledgements

We thank Christopher Potts for sharing the corpus data with us; Ashley Jin and Jure Leskovec for advice.

### References

- [1] Davison, K. P., Pennebaker, J. W., & Dickerson, S. S. (2000). Who talks? The social psychology of illness support groups. *American Psychologist*, 55(2), 205.
- [2] Coulson, N. S. (2005). Receiving social support online: an analysis of a computer-mediated support group for individuals living with irritable bowel syndrome. *CyberPsychology & Behavior*, 8(6), 580-584.
- [3] Kummervold, P. E., Gammon, D., Bergvik, S., Johnsen, J. A. K., Hasvold, T., & Rosenvinge, J. H. (2002). Social support in a wired world: use of online mental health forums in Norway. *Nordic journal of psychiatry*, 56(1), 59-65.
- [4] Pennebaker, J. W. (1997). Writing about emotional experiences as a therapeutic process. *Psychological Science*, 8(3), 162-166.
- [5] Cohen, S., & Wills, T. A. (1985). Stress, social support, and the buffering hypothesis. *Psychological bulletin*, 98(2), 310.
- [6] Uchino, B. N., Cacioppo, J. T., & Kiecolt-Glaser, J. K. (1996). The relationship between social support and physiological processes: a review with emphasis on underlying mechanisms and implications for health. *Psychological bulletin*, 119(3), 488.
- [7] Coyne, J. C., & Downey, G. (1991). Social factors and psychopathology: Stress, social support, and coping processes. *Annual review of psychology*, 42(1), 401-425.
- [8] House, J. S., Umberson, D., & Landis, K. R. (1988). Structures and processes of social support. *Annual review of sociology*, 293-318.
- [9] Miller, M., Sathi, C., Wiesenthal, D., Leskovec, J., & Potts, C. (2011, May). Sentiment Flow Through Hyperlink Networks. In ICWSM 2011.
- [10] Zafarani, R., Cole, W. D., & Liu, H. (2010). Sentiment propagation in social networks: A case study in LiveJournal. In *Advances in Social Computing* (pp. 413-420). Springer Berlin Heidelberg.
- [11] Bollen, J., Goncalves, B., Ruan, G., & Mao, H. (2011). Happiness is assortative in online social networks. *Artificial life*, 17(3), 237-251.
- [12] Potts, C. G. (2010). On the negativity of negation. In Nan Li and David Lutz, eds., *Proceedings of Semantics and Linguistic Theory 20*, 636-659. Ithaca, NY: CLC Publications.
- [13] Socher, Richard; Jeffrey Pennington; Eric H. Huang; Andrew Y. Ng, and Christopher D. Manning. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the 2011 Conference*

on Empirical Methods in Natural Language Processing, 151-161. ACL.

- [14] Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of small-world networks. *nature*, 393(6684), 440-442.
- [15] Blondel, V. D., Guillaume, J. L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), P10008.
- [16] Nielsen, F. Å. (2011). A new ANEW: Evaluation of a word list for sentiment analysis in microblogs. *arXiv preprint arXiv:1103.2903*.
- [17] Bradley, M. M., & Lang, P. J. (1999). Affective norms for English words (ANEW): Instruction manual and affective ratings (pp. 1-45). Technical Report C-1, The Center for Research in Psychophysiology, University of Florida.
- [18] Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). *Linguistic inquiry and word count: LIWC 2001*. Mahway: Lawrence Erlbaum Associates, 71.
- [19] Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, 54(1), 547-577.
- [20] Bastian M., Heymann S., & Jacomy M. (2009). Gephi: an open source software for exploring and manipulating networks. *International AAAI Conference on Weblogs and Social Media*.

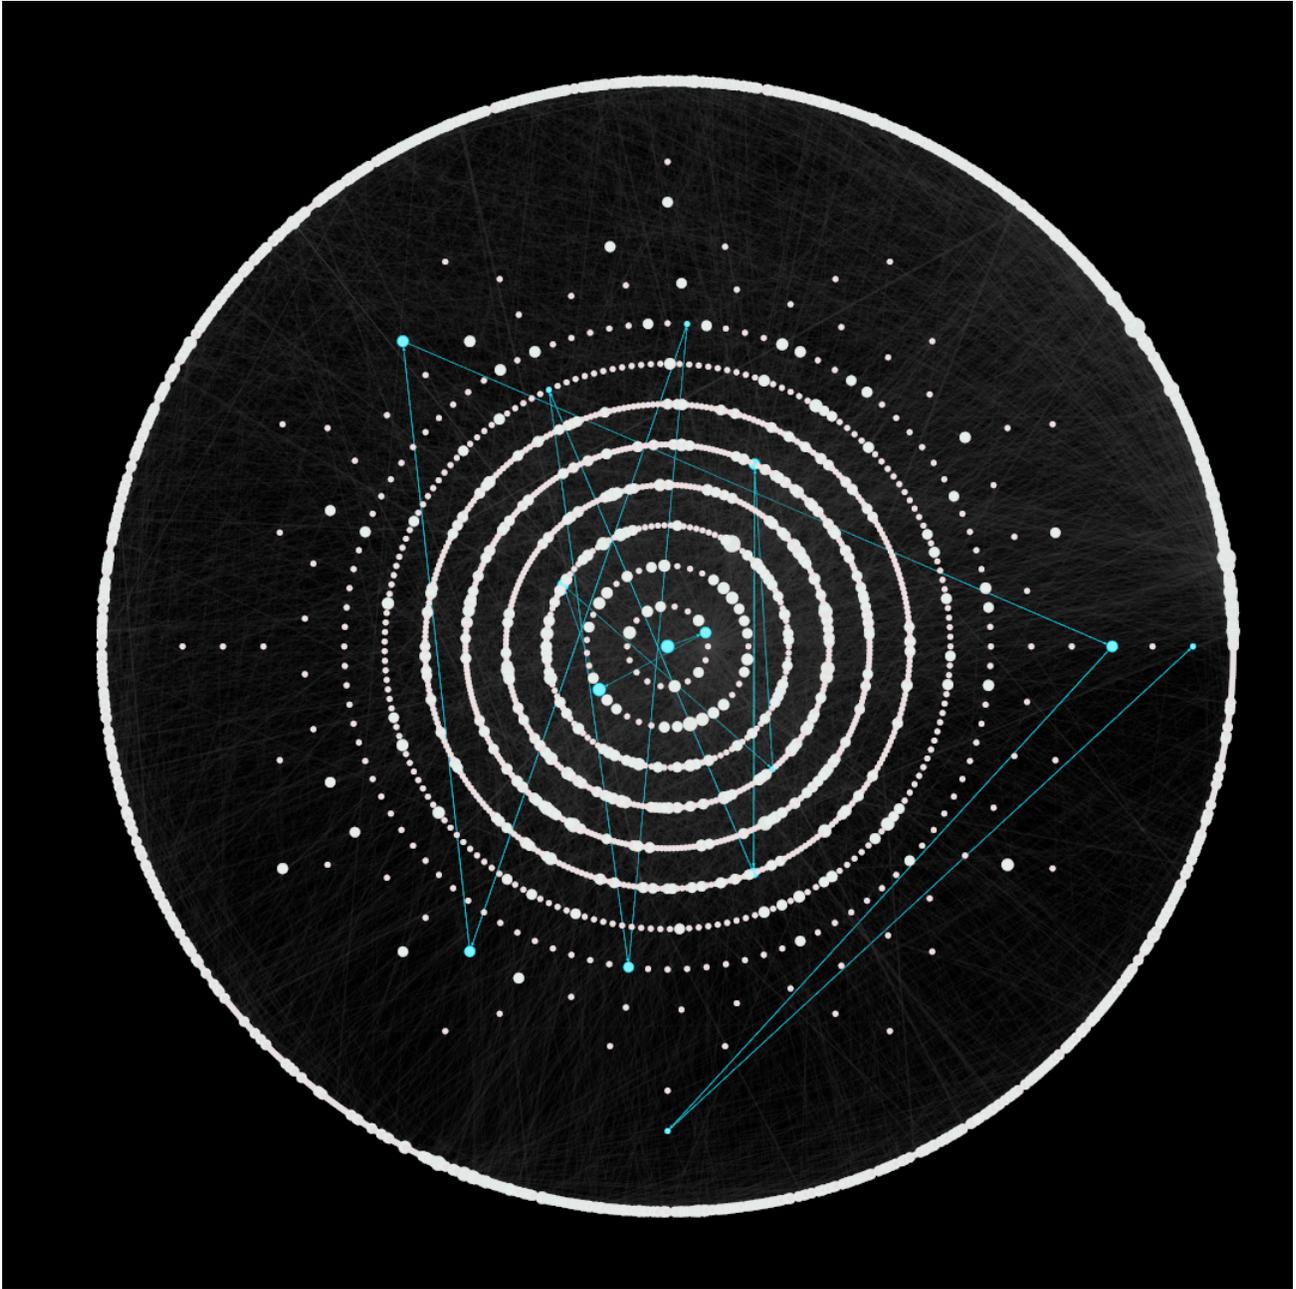


Figure 3: **Top Active Node.** Visualization of an example cascade starting from the top active node (in the center), to the node in the center of the right side of the figure. Figure was created using [20]