Using word embeddings to assess coherence within co-sharing communities of Facebook pages

networks, communities, similarity, embeddings, cyberspace

Extended Abstract

The study of network structures to investigate the emergence of polarization on online communication sites is an active research area for computational social science. However, a detailed understanding of the social dynamics undergoing the observed network is paramount to properly understand the problem at hand. In the case of social media based interaction it is often the case that, as it has been claimed by Vasques Filho & O'Neale (2020), every one-mode network should be regarded as a projection of a underlying bipartite structure (e.g. users interacting "with" a message). In this context actors can form densely connected structures (e.g. communities) for a variety of reasons that might or might not follow an homophilic principle: e.g. users can interact with a message to support it, to criticize it or just because the message was previously shared by a high-visibility user. This can easily produce noisy data that can be only partially addressed with a growing number of backboning methods (Coscia & Rossi 2019). Nevertheless, due to the lack of users' level information available for social media data, researchers often rely on methods that assume homophilic network structures to infer data about the users. Notable examples are the methods to identify the political preferences of Twitter users leveraging their following/followers or their retweet activity (Barberà 2015, Giglietto et al. 2019).

In this abstract we contribute to exploring this possible problem by proposing to use the word embedding of textual description that social media users provide of themselves to define their "semantic proximity". Moreover, we show how this can be used to integrate our understanding of networks created through online shared behaviour. By doing so we present an approach to characterize the network structure of social media entities according to how much they are semantically similar and to assess the presence (or absence) of homophily as a driving force of the observed network structure.

Method and data

To illustrate the process and the possible applications we will use a Facebook network of public pages that have been active in the public debate around climate change. The network has been constructed by manually selecting a group of actors and counter-actors (pages that oppose climate activism or even deny the concept of human-caused climate change) (N=75). These seed-actors have been monitored for 30 days for the links shared on Facebook. Crowdtangle's APIs have been then used to obtain all the other facebook pages that have shared the same links and to create a bipartite network (page-link). The projected network counts 800 nodes (pages) and 6005 weighted edges (co-sharing).

Starting from this network, to define a semantic-proximity score based on word embedding we proceeded as follow:

a) we used the Crowdtangle APIs[1] to obtain the "page description" for each page;

[1] https://github.com/CrowdTangle/API

b)we used Google's translate API[2]s to translate the page description into English;
c)we used OpenAI's ada 002 model[3] to obtain an embedding on 1536 dimensions
d)we used the R implementation of the T-Distributed Stochastic Neighbor Embedding[4] (Maaten & Hinton 2008) to reduce the data to one dimension.
e)we assigned each of the resulting semantic-proximity scores (the embedding score of the page) to the corresponding node.

This method allows us to represent the semantic diversity (or similarity) present in the page descriptions as a numeric value that we can then use in a variety of network analysis tasks.

Initial sanity-check and possible applications
Figure 1 shows the kernel density estimate of the distribution of the scores in the network. The semantic-proximity scores seem to follow a somehow multimodal distribution. Nevertheless, the network reports an assortativity value of 0.69 suggesting a positively assortative network where nodes with similar semantic proximity scores are connected. On the opposite side the network is neither assortative or disassortative with regard to the degree (0.04). Figure 2 shows a visual representation of the network where nodes have been coloured according to their semantic proximity scores. Upon manual inspection, the assortative nature of the network with pages connected with like-minded pages appears confirmed. Nevertheless, it is interesting to observe how, when combining the semantic-proximity scores with the community structure (detected with the Louvain algorithm) we can expand our understanding of the co-sharing communities. Figure 3 shows a ridgeline visualization of the distribution of the semantic-proximity scores within communities with more than 40 members. This allows us to see how there are communities that are characterized by highly homogeneous semantic-proximity while others are characterized by a larger semantic diversity. This works as a reminder that, when working with social media data, co-activity might or might not be a good proxy for homophily and that our proposed method is a good way to keep track of this phenomenon.

References

Barberá, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. Political analysis, 23(1), 76-91.

Coscia, M., & Rossi, L. (2019, August). The impact of projection and backboning on network topologies. In Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (pp. 286-293).

Giglietto, F., Righetti, N., Marino, G., & Rossi, L. (2019). Multi-party media partisanship attention score. estimating partisan attention of news media sources using twitter data in the lead-up to 2018 Italian election. Comunicazione politica, 20(1), 85-108.

Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. Journal of machine learning research, 9(11).

Vasques Filho, D., & O'Neale, D. R. (2020). Transitivity and degree assortativity explained: The bipartite structure of social networks. Physical Review E, 101(5), 052305.

[2] https://cloud.google.com/translate/docs/apis
[3] https://platform.openai.com/docs/guides/embeddings/what-are-embeddings
[4] https://github.com/jdonaldson/rtsne/

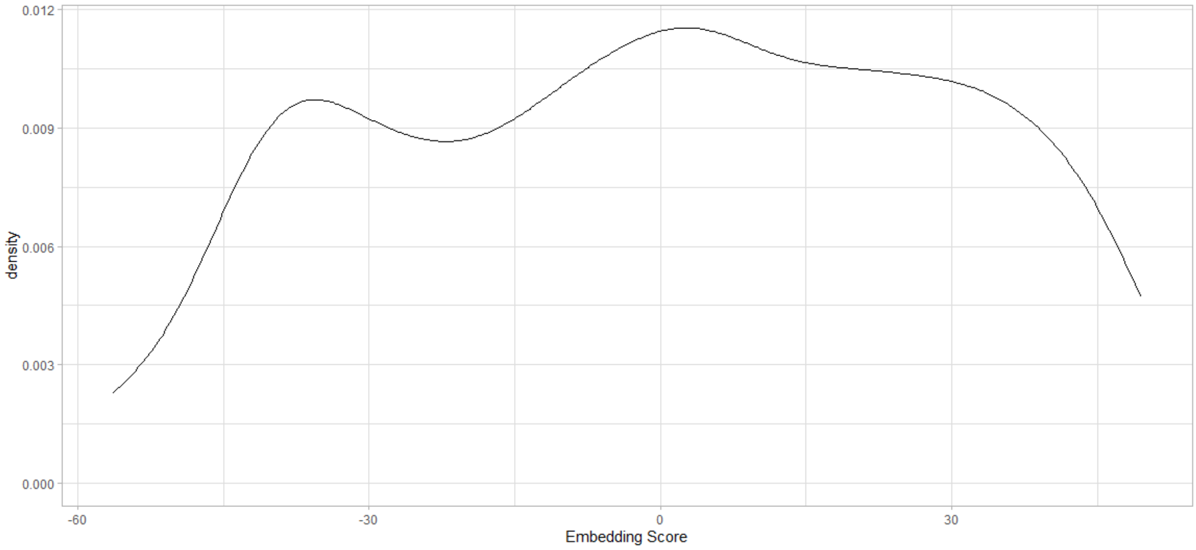

Figure 1 : Kernel density estimate of the semantic-proximity scores in the network

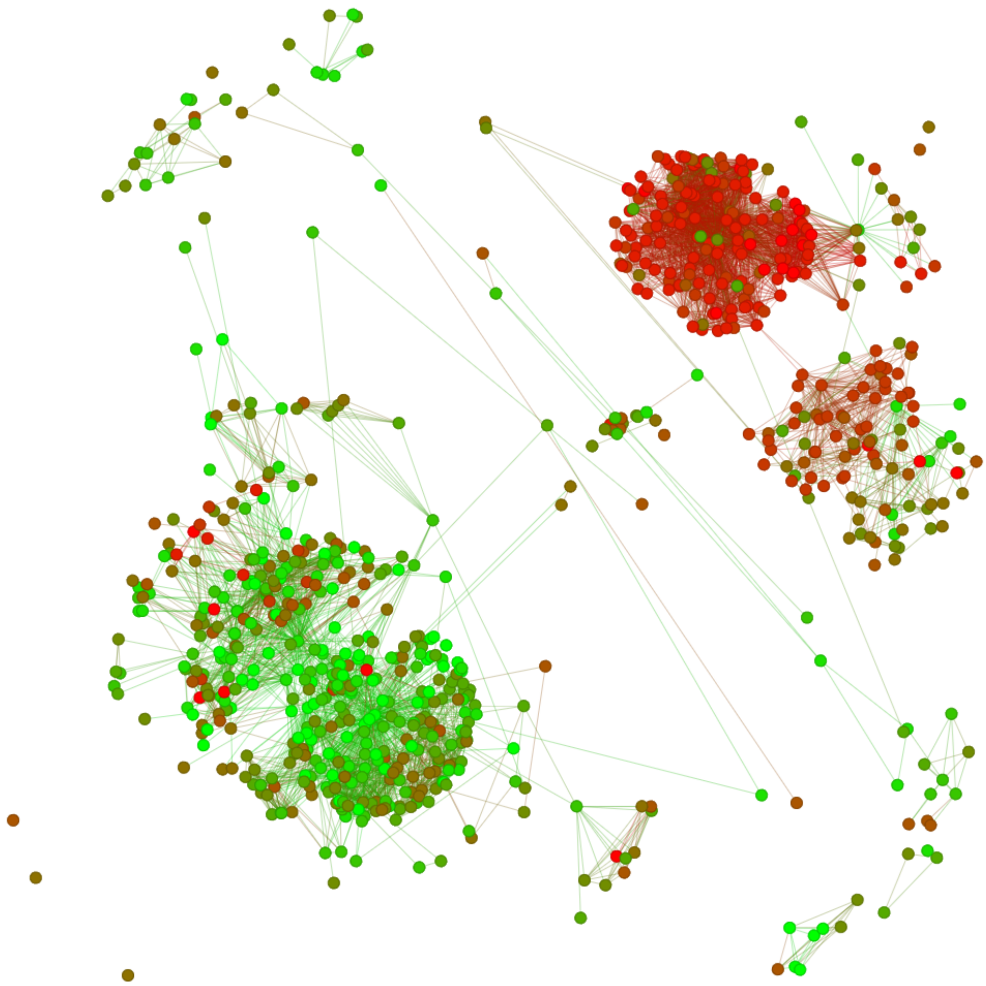

Figure 2: Network with semantic-proximity scores as color of the nodes

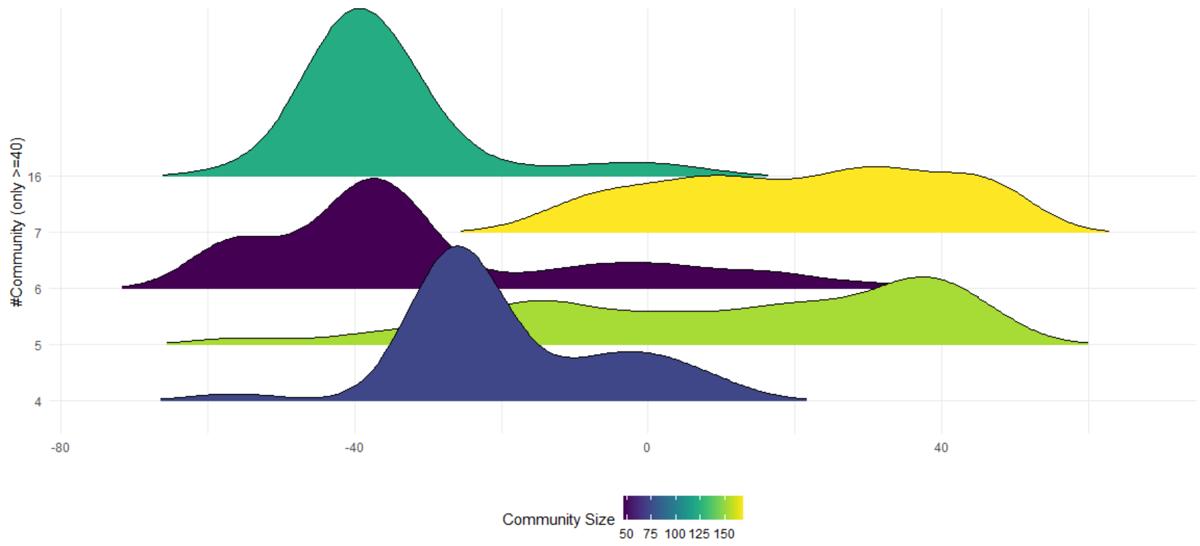

Figure 3: Distribution of the semantic-proximity scores within communities with more than 40 members

…