Community evolution with Ensemble Louvain

Bojan Evkoski[1,2], Igor Mozetič[1], and Petra Kralj Novak[1]

[1] Jozef Stefan Institute, Ljubljana, Slovenia,
[2] Jozef Stefan International Postgraduate School, Ljubljana, Slovenia,
{Bojan.Evkoski, Igor.Mozetic, Petra.Kralj.Novak}@ijs.si

A key factor in tracing communities over time is to be able to consistently detect the communities in one single time period. If one would use Louvain [1] for tracing communities over time, one would not be able to disentangle differences arising by the change in the community structure from the volatility of Louvain runs. In this paper, we present a continuation of the work, aiming to answer why Ensemble Louvain [2] is more suitable for tracking community evolution than the standard widely-used Louvain.

Ensemble Louvain [2], is a simple approach which begins by running the Louvain several times (hundred by default). Then, it builds a meta-network where a pair of nodes is connected if their total co-membership across all runs is above a given threshold (90% by default). This process forms disjoint sets of connected components in the meta-network. These sets represent our final communities. Compared to Louvain, Ensemble Louvain drastically improves stability and reproducibility of results.

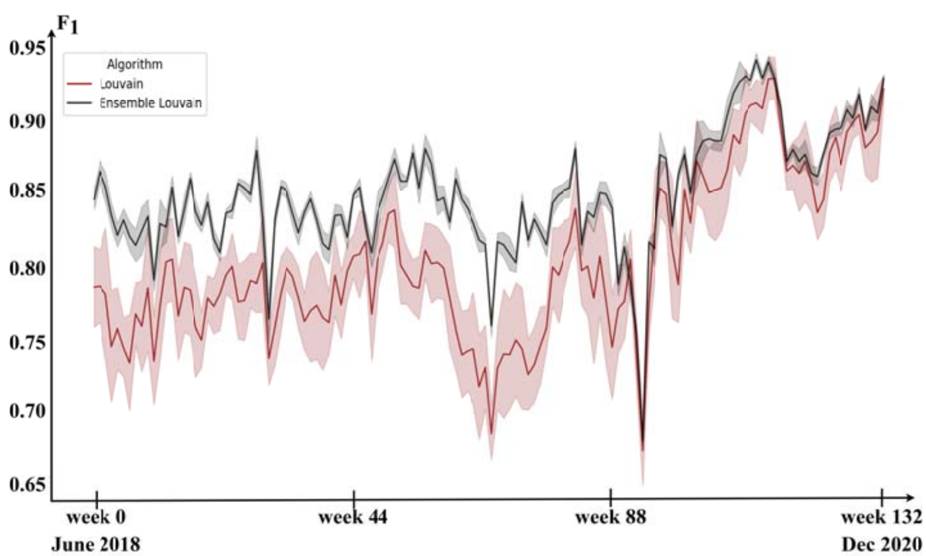

Fig. 1. Tracking community evolution — Louvain vs. Ensemble Louvain. Comparison results of the community evolution process using the two algorithms. The x-axis shows the timeline of weekly increments for which the six-month retweet networks are created. The y-axis shows the F_1 similarity score between two adjacent partitions. A lower F_1 value suggests that the community structure changed more, while higher F_1 means that it stayed more similar. An F_1 score of 1.00 means that there is no change in the network in terms of its community structure.

Community evolution is the intersection of temporal analysis and community discovery. Here, one tries to detect change in the collective behavior of groups. The problem of community evolution opens many challenges as both "parent" disciplines require different environments in order to be applicable. Temporal analysis assumes atomic granularity of events in the (dynamic) networks, while community detection requires aggregated (static) networks.

In consequence, although there is a plethora of community detection approaches (with all their advantages and drawbacks), most of the methods suffer from the common issue of unstable results (partitions) [4] [5]. This instability is due to the fact that most community detection algorithms are based on a greedy optimization of partitions on some metric (such as modularity), making them prone to produce different results each run, as they get stuck in local maxima. Accordingly, it becomes problematic in community evolution, since one cannot be sure if a change happened in the community structure, or it is a random false event caused simply by the instability of the method.

Techniques which tackle the instability issue in dynamic community detection are referred to as temporal smoothing. Adopting a specific smoothing strategy can lead to computational constraints, but it is a crucial step if one aims for finer results [3]. This is where Ensemble Louvain comes into play, applying temporal smoothing by bootstrapping, significantly stabilizing the results, and with that, removing large portion of the signal noise in the tracking of community evolution. Using ensembles as a smoothing strategy was mentioned as a prospect in the past as well [6] [7], yet without further quantitative experimental analysis of it being used. Here, we compare Ensemble Louvain with the standard Louvain and explore the impact of using the first.

As a continuation of our work [2], we apply the community evolution analysis on the Slovenian Twitter data from January 2018 to December 2020. We create retweet networks out of 24-week data, with a sliding window of one week, resulting in 133 networks, with 01.01.2018-18.06.2018 (*week 0*) being the first, and 13.07.2020-28.12.2020 (*week 132*) the last network in the sequence. Additionally, each network is created with exponential decay on the edge weights (from latest to oldest retweets), so that we prevent detecting "changes" due to lost structure patterns of the trailing data. With that, we instead emphasize the actual community behaviour shifts due to new events.

To detect community structure changes, we compare the similarity of two adjacent partitions (e.g., the *week 0* with the *week 1 network community partition*). We calculate the similarity using the F_1 metric defined in our previous work [2], which is a node-wise alternative to the common NMI and ARI scores, but works with non-identical sets as well. To compare Ensemble Louvain to the standard Louvain, we apply the described procedure ten times for the whole timeline, for both algorithms. A graphical representation of the results is shown in Fig. 1, while the statistics regarding the comparison between Louvain and Ensemble Louvain are presented in Table 1.

Fig. 1 shows that the Ensemble Louvain produces less noise (more stable results) when comparing the outputs of the ten different experiments. In terms of numbers, the average and total standard deviation (noise) of F_1 is five-fold lower for Ensemble Louvain. In practice, this means that the randomness of the process is no longer a strong factor which influences the community evolution analysis, ensuring one that the observed changes in the partitions are actually data-driven events.

Algorithm	Avg. F_1 std.	Avg. F_1 mean	Avg. F_1 coeff. of var.
Louvain	0.024 ± 0.001	0.812 ± 0.007	0.051
Ens. Louvain	$\mathbf{0.006 \pm 0.001}$	$\mathbf{0.847 \pm 0.005}$	**0.036**

Table 1. Statistics of the results visualized in Fig. 1. The "average F_1 std." shows the average y-axis standard deviation for the ten experiments across the whole timeline. Lower deviation allows better interpretability of the community evolution. The "average F_1 mean" is the average y-axis value, considering the darker line (mean of ten experiments). Higher F_1 suggests generally more similar (less volatile) partitions through the y-axis. Finally, the "average F_1 coefficient of variation", shows the average volatility of the ten experiments across the x-axis.

We also observe a generally higher F_1 score throughout the whole timeline for the Ensemble Louvain. This means that, according to the Ensemble Louvain outputs, the adjacent partition differences are significantly lower compared to when analyzed using Louvain. In other words, the evolution is less volatile. This is also confirmed by the calculated coefficient of variation. One of the effects of using Ensemble Louvain is, that if a node does not firmly belong to one particular group (i.e., shifts membership in the single Louvain runs), it will be left out as an "outsider", not affiliating with any of the communities. Most of these nodes do not change their behaviour rapidly, so they usually maintain their non-affiliation, increasing the general F_1 score of the adjacent networks. Additionally, having these unstable nodes out of the large communities, the similarity between these communities becomes even higher, ending with a consistently higher score compared to the standard Louvain. Finally, tracking community evolution benefits from these behaviours, as the noisy variables are removed from the process of detecting change, making Ensemble Louvain preferred over the standard Louvain.

Acknowledgments. The authors acknowledge financial support from the Slovenian Research Agency (research core funding no. P2-103), and the European Union's Rights, Equality and Citizenship Programme (2014-2020) project IMSyPP (grant no. 875263). The content of this abstract represents the views of the authors only and is their sole responsibility. The European Commission does not accept any responsibility for use that may be made of the information it contains.

References

1. V. D. Blondel, J. Guillaume, R. Lambiotte and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, p.P10008, 2008(10)
2. B. Evkoski, I. Mozetič, N. Ljubešić and P. K. Novak. Community evolution in retweet networks. *PLoS ONE*, 16(9):e0256175, 2021. *Non-anonymized version available from: https://arxiv.org/abs/2105.06214*
3. G. Rossetti and R. Cazabet. Community discovery in dynamic networks: a survey. *ACM Computing Surveys (CSUR)* 51.2(1-37), 2018.
4. T. Aynaud and J. Guillaume. Static community detection algorithms for evolving networks. *8th International symposium on modeling and optimization in mobile, Ad Hoc, and wireless networks*. IEEE, 2010.
5. Bommarito II, J. Michael, D. M. Katz and J. L. Zelner. On the stability of community detection algorithms on longitudinal citation data. *Procedia-Social and Behavioral Sciences 4*. 26-37, 2010
6. T. Chakraborty, P. Noseong, A. Aysh and V.S. Subrahmain. Ensemble detection and analysis of communities in complex networks. *ACM Transcations on Data Science 1.1*. 1-34, 2020
7. V. Poulin and F. Théberge. Ensemble clustering for graphs. *International Conference on Complex Networks and their Applications* 231-243, 2018