Combining Network and Language Indicators for Tracking Conflict Intensity

Anna Rumshisky[⋆*], Mikhail Gronas[†], Peter Potash[⋆], Mikhail Dubov[‡], Alexey Romanov[⋆], Saurabh Kulshreshtha[⋆], Alex Gribov[⋆]

[⋆]Department of Computer Science, University of Massachusetts Lowell
[†]Department of Russian, Dartmouth College
[‡]Higher School of Economics, Moscow, Russia
[*]contact: arum@cs.uml.edu

Abstract. This work seeks to analyze the dynamics of social or political conflict as it develops over time, using a combination of network-based and language-based measures of conflict intensity derived from social media data. Specifically, we look at the random-walk based measure of graph polarization, text-based sentiment analysis, and the corresponding shift in word meaning and use by the opposing sides. We analyze the interplay of these views of conflict using the Ukraine-Russian Maidan crisis as a case study.

1 Introduction

Over the past decade, social media websites such as Twitter and Facebook (as well as their counterparts in other countries, such as Weibo in China or VKontakte in Russia) have become an integral part of social life in many locations around the world. In conjunction with the proliferation of social media use, a number of political conflicts and incidents of social unrest occurred around the world (for example, events associated with the "Arab spring" in Egypt and Tunisia in 2010 and 2011, Spanish 15-M movement of 2011-2012, and the Russian takeover of Crimea in 2014). As a result, substantial amounts of data have been accumulated regarding people's behavior (both linguistic and extra-linguistic) in social media prior to and during such conflicts. We are now in a position to attempt a data-driven approach to conflict dynamics, detect internal logic, and analyze patterns in how conflicts originate and develop. In doing so, we follow in the footsteps of conflict sociologists such as Randal Collins who in his seminal 2012 paper outlined the dynamic processes involved in the initiation, development, and de-escalation of conflict [2].

In this paper, our goal is to look at the conflict dynamically as it develops. We hypothesize that as a complex phenomenon, conflict gets reflected in multiple related processes. Because conflict in our context is an inherently intra-human phenomenon, signals of its dynamics are likely present in common forms of communication, of which social media has become ubiquitous. If that is indeed the case, multiple conflict indicators derived from social media can be combined into a composite measure of conflict intensity. We view conflict as a systemic and dynamic phenomenon, in which these indicators can be tracked over time to detect significant rapid changes. We argue that a typology of interaction dynamics for different indicators may be used to characterize the development of conflict over time.

For a composite measure to work, there should be a meaningful relationship between different components of a composite measure. In this paper, we investigate whether there is such a relationship between language-based and network-based measures of controversy. While previous work such as Garimella et al. [6] has substantially advanced the study of social media-based controversy, by providing the first empirical validation for the effectiveness (or ineffectiveness) of various network-based and language-based controversy measures, there has been no work so far that seeks to determine whether such controversy measures have a meaningful temporal relationship. In order to analyze the interplay between these two complementary views of conflict, we use the 2013-2014 Ukraine-Russia Maidan crisis as a case study. By having data collected over a year-long time period, we are able to calculate controversy measures at various intervals as the Maidan crisis evolves. This case study is based on the Russian-speaking social media during the Ukrainian events of 2014, where civil discontent and division led to protests in the streets, and eventually to armed violence.

We examine the network-based controversy measure using the user graph in which connections are induced by information-sharing/information-consuming patterns in the network. We build and examine network graphs in which an edge connects two users if they have liked the same social media post. For the language-based measures, we analyze the rhetorical patterns both quantitatively and qualitatively: the overall sentiment expressed by the opposing sides (determined using automated text-based sentiment analysis), and the shift in word meaning and use ("lexical drift") that happens, as the language used by opposing communities to describe events related to the conflict develops diverging semantic representations. We analyze the drift in word meaning using a word embedding model [13] that uses contexts of a word's occurrences in order to create dense vector space embeddings for individual words.

2 Related Work

The study of conflict and civil unrest is a highly interdisciplinary field at the intersection of philosophy, psychology, economics, sociology and political science, to name just a few. In the context of our study, especially relevant are big data oriented political science projects such as GDELT (Global Database of Events, Language, and Tone) and ICEWS (Integrated Crisis Early Warning System) that aim to predict and monitor civil conflicts on the basis of large-scale analysis of news sources. Distinct from such approaches, we focus on a social media representation of conflicts and network properties in the conflicting communities. Our approach focuses on the micro-analysis of the underlying social dynamics that lead to conflict, operating at the level of individual political opinions and allegiance.

In the field of theoretical political science and social psychology, Nowak et al [17] and Deutsch et al [3] have previously suggested that one of the defining properties of conflict is the shift and eventual alignment of the opposing opinions across the divisions between social groups. While our approach is based on similar underlying ideas, our model allows us to determine empirically and track such dynamic changes as they develop. Alternatively, the field of network science conceptualizes conflict in terms of polarization between user communities. Over the past few years, the problem of community detection in social networks has received a lot of attention. In 2006, Newman [16] defined the modularity metric Q that detects how separate two communities are from each other. Specifically, Q examines the clusters formed by two distinct community groups in a network, and how this clustering compares to a random network. The higher its value, the more modular the network is. Blondel et al. [1] proposed a community detection algorithm that tries to maximize Q in a given network. Furthermore, the algorithm generates a hierarchy of partitions: for each partition it makes, that partition has its own subpartitions. This view becomes more microscopic until each node is in its own community.

Peixoto [18] states that the metric of modularity does not take into account the possible statistical fluctuations of the null model, and notes that modularity can detect highly modular communities in random graphs [8]. Furthermore, Blondel et al.'s algorithm has difficulty detecting communities when the cluster sizes are small [5, 12]. Another criticism of using the modularity metric for polarized community detection comes from the recent work by Guerra et al [7], who note that modularity and related measures recover separate and autonomous communities, but not opposing user clusters. They propose a measure based on the notion of *popularity-at-boundary* in order to capture the polarization reflecting the opposition between user clusters.

There are several advantages and disadvantages to the graph-based measures proposed for quantifying controversy/polarization [6, 7, 15]. These measures analyze the interaction between to two communities in a graph, based on differing theories of how conflict/polarization manifests itself. Garimella et al. [6] have proposed the best experimental setup to quantify the best controversy measure. The authors do so by selecting both controversial and non-controversial topics, and show how well a given measure does at differentiating between the two types of topics. The authors show that a conflict measure based on random walks between the two communities (RWC) performs the best out of five methods tested [6, 7, 15]. Furthermore, the authors show that the standard deviation of sentiment is a strong linguistic marker of controversy, whereas average sentiment and divergence of the conflicting communities' vocabularies are not able to separate between conflicting and non-conflicting topics.

Volkova et al [21] is another example of linguistic analysis during controversy. The authors used language analysis to study and predict the emotional response during the Maidan crisis. In contrast to their approach, we do not rely on noisy location data to create our corpus, but instead use self-labeled user groups relevant to the crisis.

3 Dataset

In this study, we used data from Russian-speaking online media, posted during the Ukrainian events of 2013-2014. We use the largest Russian social network "VKontakte" (VK)[1]. According to liveinternet.ru, VKontakte has 320 million registered users and is the most popular social network in both Russia and Ukraine. During the conflict, both pro-Russian and pro-Ukrainian side (also known as "Antimaidan" and "Pro-" or "Evromaidan") were represented online by large numbers of Russian-speaking users.

We have built a scalable open stack system for data collection from VKontakte using VK API. The system is implemented in Python using a PostgreSQL database and Redis-based message queue. VK API has a less restrictive policy than Facebook's API, making it an especially suitable social network for research. Our system supports such API methods as retrieving group members, retrieving all posts from a wall, retrieving comments and likes for a given post, and so on. Moreover, we are able to collect all posts (and its related attributes) by a user that are public.

In order to seed the data collection, we selected the most popular user groups from the two opposing camps, the Evromaidan group (154,589 members) and the Antimaidan group (580,672 members). We then manually annotated other groups to which the administrators of these two groups belonged, selecting the groups with political content. This process produced 47 Evromaidan-related groups with 2,445,661 unique members and 51 Antimaidan-related groups with 1,942,918 unique members. We retrieved all posts from these group walls, as well as all the users who have liked one of the posts.

We then selected the users who were reasonably active during the time period of interest. We defined *active users* as those who averaged 2 or more posts per 3 months at least once over the target time frame (Oct 1, 2013 - Oct 1, 2014). This resulted in 745,880 users from the Antimaidan-related groups and 725,053 users from the Evromaidan-related groups. We retrieved all posts from these user walls, as well as all the users who have liked these posts. In order to restrict our data to politically-themed postings, we built a list of 45 political keywords, which included the names of political figures, locations, and derogatory terms used by both sides. This list was used to filter the wall posts of Evromaidan and Antimaidan users.

4 Methodology

Our work leverages the complementary views of networks and language when analyzing controversy. The first part of this section describes the quantitative

[1] `http://vk.com`

methods that we use use for measuring conflict. Because these methods are general and provide a single numeric value, they have the potential to be combined into a composite index. Next, we discuss our qualitative language-based method of lexical drift. We finish by discussing how to generalize our methodology to the analysis of other events.

4.1 Network-Based Controversy

Our network-based controversy measure is the RWC measure from Garimella et al. [6]. For all graphs, nodes represent users. We create an edge between two users (nodes) if they have liked the same post. We use the python package NetworkX[2] to construct and manipulate graphs. Once we construct a graph, we extract the largest connected component. The RWC measure assumes that a given graph has two communities already identified − the graph has two clusters for the nodes. For the VK data, based on the data collection methodology, we have predefined communities based on whether the user came from an Antimaidan or Evromaidan group (see Section 3).

To compute the RWC measure on a graph, the first step is to identify the k nodes with highest degree in each community. These are referred to as the authoritative nodes. Generally speaking, the goal is to calculate the probabilities of starting a walk at a random node in a given community, and end at an authoritative node in same community, as well as ending at an authoritative node in the opposite community. In practice, these probabilities are calculated through inference as follows: (1) Randomly select 10% of nodes in each community, (2) For each node, perform a random walk until an authoritative node is reached in either community. We repeat steps 1 and 2 a thousand times to best estimate the target probabilities. For each random walk, we keep track of the counts related to starting on a given side and ending on a given side. The final formula for RWC is:

$$RWC = P_{XX}P_{YY} - P_{XY}P_{YX} \tag{1}$$

where

$$P_{AB} = P[start\,in\,community\,A|\,end\,in\,community\,B] \tag{2}$$

This measure will produce a real-valued score for a given graph, which in turn represents actions of users in a network.

4.2 Language-Based Controversy

The quantitative linguistic attribute we measure is the sentiment of language used in the posts from the VK data. Sentiment analysis produces a discrete label for a piece of text: positive, negative, or neutral, which we convert to an integer using the following mapping: positive → 1, neutral → 0, negative → -1. Since these values are given per piece of text (post), we compute aggregate statistics: average (mean) sentiment and standard deviation of sentiment. Garimella et al.

[2] https://networkx.github.io/

[6] posit that the standard deviation of sentiment is a conflict indicator, like RWC. We apply language-based analysis to the data used to construct the user graphs. Thus, the exact posts we analyze is a subset of the overall data that has been collected.

Sentiment Analysis in Russian In order to measure the sentiment for Russian, we used a Python port of the Sentimental system[3] for sentiment analysis. Sentimental is a dictionary-based sentiment analysis system with basic capabilities of handling negated words. We used a lexicon which consisted of 7640 words with sentiment scores from -5 to 5. The overall sentiment score of a sentence is a sum of the scores of individual words, normalized be the length of the sentence. This system, along with the used lexicon, is publicly available on GitHub[4]. We threshold the output of the system to predict its label as follows: $> 0 \rightarrow positive$, $= 0 \rightarrow neutral$, $< 0 \rightarrow negative$.

4.3 Quantitative Temporal Analysis

In order to analyze the temporal trends of our case studies, we segment the available data into time slices, which generally is dictated by what data is at our disposal. Based on the non-overlapping data segments, we calculate the following measures at each time slice: RWC, average sentiment, and standard deviation of sentiment. For the Maidan case study, the time slices are at monthly intervals. Between group and user wall posts, there is an average of 121,989 posts per time slice.

4.4 Qualitative Linguistic Analysis: Lexical Drift

Following the distributional hypothesis [9, 14, 4], the contexts in which words occur are indicative of their meanings. We trained word2vec embeddings [13] for each temporal slice of VK data, following the methodology outlined in [11]. A similar methodology has previously been used to analyze language during conflict [20]. Word2vec creates a mapping between words and vectors in $\mathbb{R}^n$, where n is fixed. Moreover, Mikolov et al. argue that the topology of the vector space is also *semantically* continuous – geometric proximity equates to semantic similarity. As new data is available at each time slice, allowing the model to continue training, the positioning of word representations in the vector space shifts as words begin to appear in different contexts. To compute distance between embeddings e_i, e_j, we use the following formula:

$$CosineDistance(e_i, e_j) = 1 - \cos(\theta_{e_i,e_j}) \tag{3}$$

where $\cos(\theta_{e_i,e_j})$ is the cosine of the angle between the embedding vectors e_i, e_j[5], which is in the range $[-1, 1]$, equaling -1 when the vectors point opposite directions and equaling 1 when the vectors point the same direction. Thus, a

[3] https://github.com/Wobot/Sentimental
[4] https://github.com/text-machine-lab/sentimental
[5] https://en.wikipedia.org/wiki/Cosine_similarity

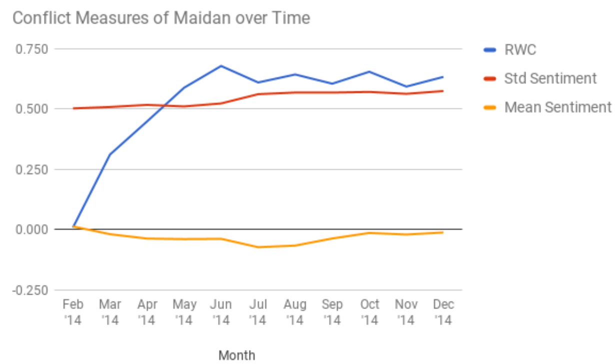

Fig. 1. Measures of conflict for Maidan, shown in monthly segments

higher cosine distance translates to the meanings being farther apart. We use the Python package Gensim[6] to train word vectors.

4.5 Generalizability of Methodology

We argue that our method is generalizable across conflicts, given the presence of social media data (collected from a meaningful time period). We do note that the users of opposing sides should be in the same network and use the same language, or at least have familiarity with the opposing side's language. While our methodology uses joint liking to create edges between users, if Twitter data were to be used, one can use retweets, as Garimella et al. [6] suggest, instead. Secondly, the RWC measure requires the presence of two predefined communities. Although our data collection process is able to annotate this directly, one could use existing algorithms such as Metis[7] [10]. Lastly, our work requires a sentiment analysis tool. However, off-the-shelf sentiment tools for languages other than English may been difficult to come across. Annotated corpora for languages other than English are beginning to appear, such as for Arabic [19]. Even if annotated sentiment data is not available, one may use a dictionary-based approach such as ours.

5 Results and Discussion

In this section, we presents and discuss the quantitative and qualitative measures of the Maidan Crisis's temporal dynamics, based both on network and linguistic

[6] `https://radimrehurek.com/gensim/`

[7] We tested the Metis algorithm on our own data and found it recorded 80% accuracy predicting community membership.

analysis. The interplay between the random walk controversy (RWC) measure and the dominating sentiment during the investigated time period is shown in Figure 1. For our case study, the relationship between these two reflections of an ongoing conflict confirmed our initial hypothesis. Essentially, as the conflict intensifies, both the RWC measure and the standard deviation of overall sentiment expressed by the opposing groups (SenSTD) will increase in unison. And in fact, we found that RWC and SenSTD are positively correlated, with Pearson and Spearman correlation values of 0.674 and 0.745, respectively. We also observed that RWC and the average of the absolute value of the overall sentiment correlate negatively (with Pearson and Spearman correlation values of -0.598 and -0.291, respectively), confirming that negative sentiment accompanies the intensification of conflict.

5.1 Information-Sharing User Network

In order to visualize the evolving conflict during the Maidan crisis in Ukraine, we created graphs based on VK data from specific temporal intervals (See Figure 2). Users are represented as nodes, and an edge exists between two users in a graph if they have liked the same post (either on user walls or on group walls). In order to induce graphs from specific time periods, we restricted a given graph to likes that occurred on posts that were created in a certain interval. We used a 12-month time period that started October 2013 and ended September 2014. In order to create the graph visualizations, we used the Python library NetworkX. The graph layout we have chosen is based on the Fruchterman-Reingold layout, which is a force-directed algorithm[8]. The algorithm simultaneously tries to minimize the distance between highly connected nodes while maximizing the distance between minimally connected nodes. Blue nodes represent pro-Maidan users, and red nodes represent anti-Maidan users.

As we proceed from the beginning of the conflict, we can observe the initial cluster formation, as a chaotic mass of users organizes into clearly defined clusters, based on the sources they like. By December 2013 (the month when support for the Maidan protest became widespread throughout Western and Central Ukraine), we see the formation of the Maidan cluster, still unopposed. Around January 2014, we observe the appearance of the counter-cluster, coinciding with the growing organization and ideological coherence of the anti-Maidan forces, primarily in the Eastern regions. As the conflict intensifies, the clusters grow more dense and modular, which corresponds to the increasing hostility between the opposing factions in real life and the flare up of the open military confrontation. Significantly, around January 2014, we also begin to observe the formation of a "bridge" between the two groups, i.e. a set of users who clearly like both groups of sources. Interestingly, these boundary, "bridge" users in this case seem to self-identify as anti-Maidan (they belong to the anti-Maidan user groups, and are correspondingly shown in red).

[8] https://en.wikipedia.org/wiki/Force-directed_graph_drawing

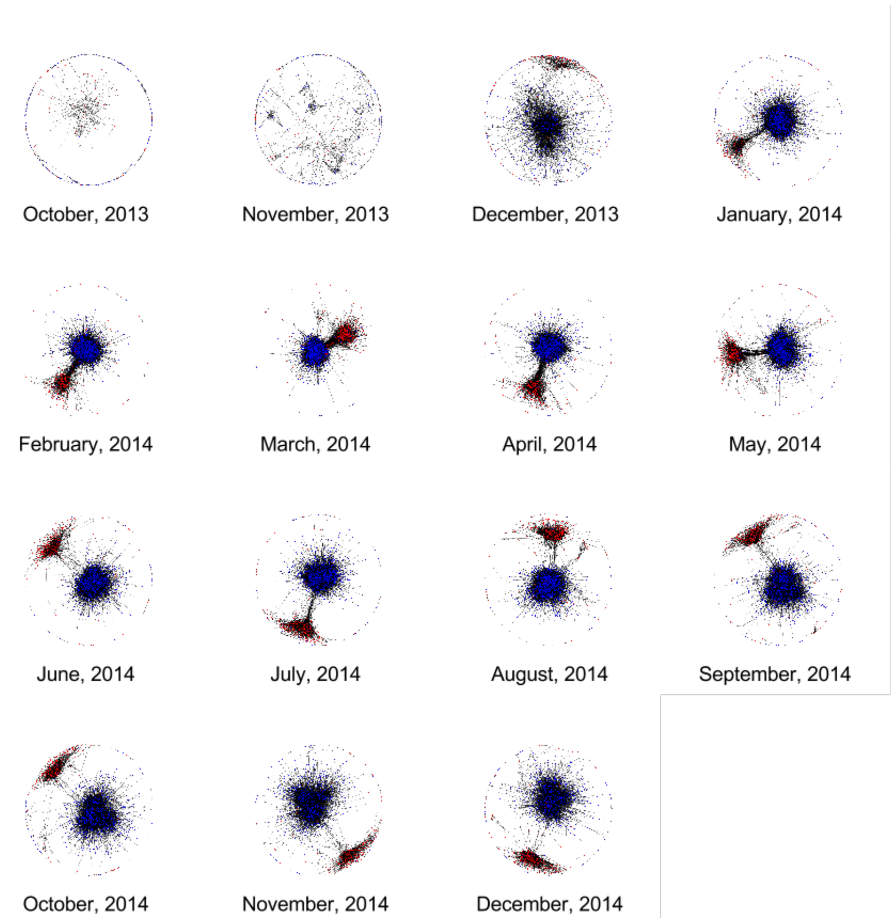

Fig. 2. Visualization of pro- and anti-Maidan VK user networks as conflict develops. Blue nodes represent Evromaidan users, and red nodes represent Antimaidan users.

Essentially, in Figure 2, you can see a clear reflection of the conflict dynamics described by Collins [2]. October and November 2013 graphs show a random arrangement of nodes in the absence of conflict, followed by an explosion that begins in December 2013 and January 2014. A bridge is formed and then disappears, as polarization increases and the conflict drives out the neutrals. The bridge begins to thin out around May 2014 and disappears completely in June 2014. As the bridge no longer exists, there is virtually no way for the representatives of the opposing sides to experience the discourse from the other side. The network visualization in this figure thus clearly shows the logic hypothesized by Collins. This corresponds to the plateau observed in Figure 1 that shows the random walk controversy measure for the Maidan data.

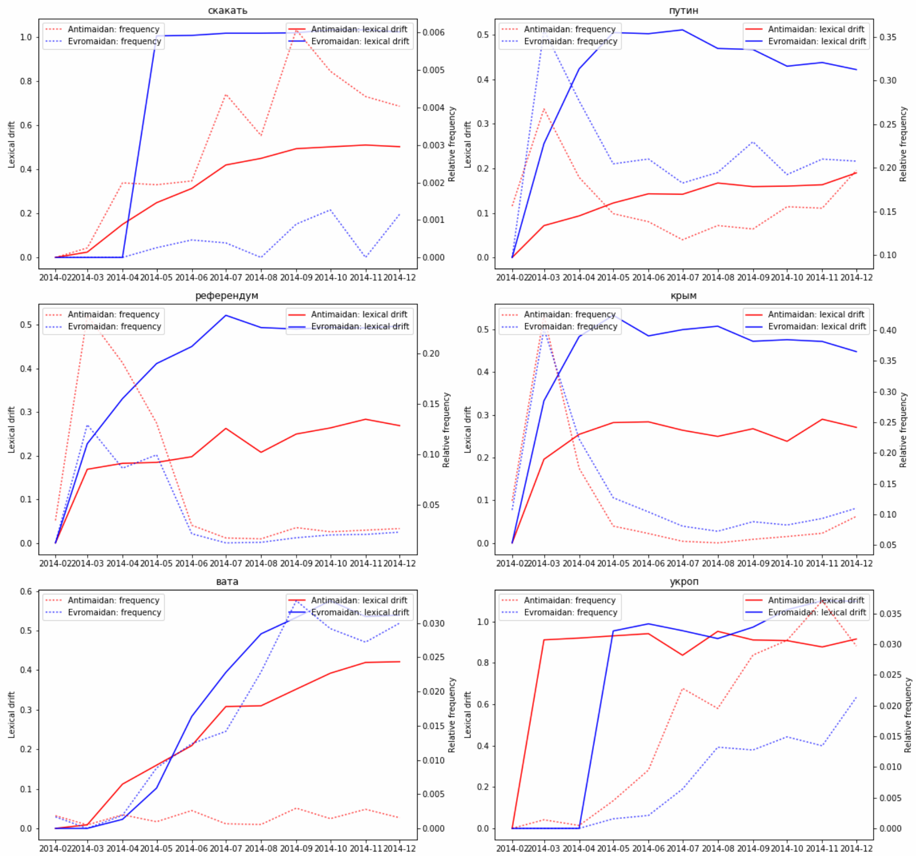

Fig. 3. Important "drifter" words that changed meaning the most during the conflict. The graph shows the cosine distance (left Y axis gives the scale) between the initial vector space embedding in Feb 2014 and the embeddings for the following months. Dotted lines show the relative frequency for each word, i.e., the percentage of posts wherein the word appears (right Y axis gives the scale). The following words are shown (left to right, top down): скакать (skakat') "jump", Путин "Putin", референдум "referendum", Крым (krim) "Crimea", вата (vata) "cottonballs", укроп (ukrop) "dill" or "Ukrainian nationalist".

5.2 Lexical Drift

The change in the vector space embedding for the same word over time allows us to track word meaning drift as it happens during the conflict. Figure 3 shows the change in contextual distribution of important "drifter" words which changed meaning the most during the conflict. The graph shows the cosine distance between the initial vector space embedding in Feb 2014 and the embeddings for the following months.

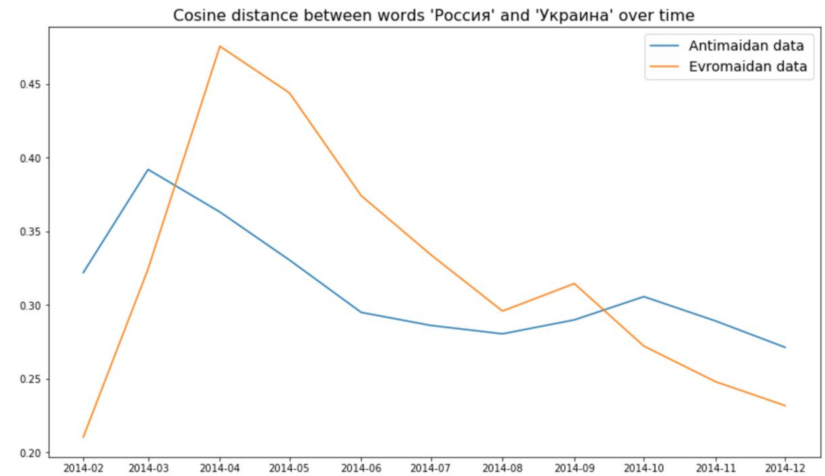

Fig. 4. Cosine distance between the words Russia and Ukraine in Evromaidan and Antimaidan data

The word “скакать”(skakat) (Figure 3, top left), for example, became a popular taunt, a derogatory term applied by the anti-maidan movement to all various types of Ukrainian nationalist activities. This use of the verb originated in the viral YouTube video that captured a crowd of young Ukrainian nationalists braving the winter cold and trying to keep warm by hopping together and chanting “кто не пляшет тот москаль”, “who is not jumping is a moscovite”. The chant became the object of ridicule among the “moscovites” and the word “скакать” (skakat) drifted towards new (political and anti-Ukrainian contexts). Cosine distance plot for “skakat” suggests that the shift away from the original meaning for this word was the most pronounced amongst the words shown.

The graph for the word “Путин” (top right) shows a drastic increase in drift for Putin’s name among the Euromaidan users from the very start of the conflict. The corresponding drifting rise on the anti-maidan side is much flatter. This is explained by the singular importance of the Russian president in the conflict. In the eyes of Evromaidan community, he became the personification of the conflict (and of the evil in general). Putin became the personal enemy of the independence minded Ukrainians and the favorite object for taunts and insults – thus drifting from political contexts to such new embedding neighbours as “хуйло” (dick) in September 2014.

The word “вата” (vata) comes from the word “ватник” (vatnik), which is a type of heavy cotton jacket popular in rural Russia and often worn by soldiers, prisoners, and farmers. The word “vata” itself means “cotton balls”. During the conflict, the term “vata” became a pejorative term picked up by the Evromaidan side as a way to taunt Russian nationalists and Russian-speaking Ukrainians (essentially the Antimaidan side), with the connotation somewhat similar to the usage of "redneck" in American English. This change in meaning is clearly visible in the corresponding lexical drift graph. Note that the term's use by the Antimaidan side has also shifted substantially (even if to a smaller relative extent compared to the Evromaidan side). Moreover, we are able to quantitatively capture the fact that this term was later re-appropriated to some extent by the Russian nationalists who started to use it as a pride badge of sorts. However, the frequency plot for "vata" shows that this takeover was still not comparable to the pervasive pejorative use by the Evromaidan side.

Укроп (ukrop) is a word meaning a garden herb, dill, which happens to start with the same combination of sound "UKR" as the word Ukrainian. This led to its appropriation by the Antimaidan side as a derogatory term for Ukrainians. Later on, the word was re-appropriated by Ukrainian nationalists, and was even used as a name for a pro-Maidan political party. These processes are illustrated by the graphs of semantic drifts: the word first shifts its position in the embedding space in the Antimaidan discourse (from dill to an anti-Ukrainian tease), and then undergoes a similar, albeit weaker, transformation among pro-maidaners. Furthermore, the up-tick in lexical drift in the Evromaidan graph that starts in September 2014 coincides with the formation of the aforementioned political party.

The words "референдум" (referendum) and "Крым" (krim; Crimea in English) see an immediate change in meaning (high positive slope) going from February to March as there were referendum votes in Crimea as well the Donetsk and Luhansk oblasts (provinces) in March that represented pro-Russian desire to secede from Ukraine and potentially join the Russian Federation. This change in meaning levels off, though, in the remaining months as the conflict continues.

This view of the lexical change also allows us to see relative drift in meaning of initially similar words. For example, consider the comparison in Figure 4 which shows that the cosine distance between the words Russia and Ukraine in Euromaidan data increased drastically in April 2014. This rise coincides with with the escalation of the conflict and the beginning of its military phase. On April 6, 2014, the Eastern separatists captured an administrative buildings in Donetsk and Lugansk regions and soon proclaimed independence. The occupation of Slavyansk by the pro-Russian forces led by Igor Strelkov began on April 12. Thus, in the Euromaidan discourse Russia became associated with the open aggression and intrusion into the Ukraine's internal affairs. This lead to a drastic peak in the semantic drift that gradually subsides in the next few months.

6 Conclusion

We suggested methods for an analysis of the temporal dynamics of the political conflict as reflected in social media, using the 2014 Russian-Ukrainian conflict as a case study. We analyzed the interplay of the division of network-based vs. language-based measures of conflict, using Random Walk Controversy as a network measure and standard deviation of sentiment and semantic drift as verbal measures. We investigated the hypothesis that the network-based and the language-based measures of conflict intensity should correlate and provided a preliminary statistical confirmation for it. We also provided a data-driven illustration for Randal Collins' influential theory of conflict development. Specifically, we showed for that the major stages of conflict as described by Collins (explosion, plateau, dissipation) and some of the behavioral patterns he postulated (including driving out of the neutrals) may be observed and visualized in the social media data for the 2014 Russian-Ukrainian conflict. Finally, we performed a qualitative analysis of the semantic shift for key terms that typified the conflict.

Acknowledgments

This work was supported in part by the U.S. Army Research Office under Grant No. W911NF-16-1-0174.

References

1. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. Journal of Statistical Mechanics: Theory and Experiment 2008(10), P10008 (2008)
2. Collins, R.: C-escalation and d-escalation: A theory of the time-dynamics of conflict. American Sociological Review 77(1), 1–20 (2012)
3. Deutsch, M., Coleman, P.T., Marcus, E.C.: The handbook of conflict resolution: Theory and practice. John Wiley & Sons (2011)
4. Firth, J.R.: A synopsis of linguistic theory 1930-1955. In: Studies in Linguistic Analysis, pp. 1–32. Oxford: Philological Society (1957)
5. Fortunato, S., Barthélemy, M.: Resolution limit in community detection. Proceedings of the National Academy of Sciences 104(1), 36–41 (2007)
6. Garimella, K., De Francisci Morales, G., Gionis, A., Mathioudakis, M.: Quantifying controversy in social media. In: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining. pp. 33–42. ACM (2016)
7. Guerra, P.H.C., Meira Jr, W., Cardie, C., Kleinberg, R.: A measure of polarization on social media networks based on community boundaries. In: ICWSM (2013)
8. Guimera, R., Sales-Pardo, M., Amaral, L.A.N.: Modularity from fluctuations in random graphs and complex networks. Physical Review E 70(2), 025101 (2004)
9. Harris, Z.: Distributional structure. In: Katz, J. (ed.) Philosophy of Linguistics, pp. 26–47. Oxford University Press, New York (1985)
10. Karypis, G., Kumar, V.: Metis–unstructured graph partitioning and sparse matrix ordering system, version 2.0 (1995)
11. Kim, Y., Chiu, Y.I., Hanaki, K., Hegde, D., Petrov, S.: Temporal analysis of language through neural language models. ACL 2014 p. 61 (2014)
12. Lancichinetti, A., Fortunato, S.: Limits of modularity maximization in community detection. Physical Review E 84(6), 066122 (2011)
13. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed representations of words and phrases and their compositionality. In: In Proceedings of NIPS (2013)
14. Miller, G., Charles, W.: Contextual correlates of semantic similarity. Language and Cognitive Processes 6(1), 1–28 (1991)

15. Morales, A., Borondo, J., Losada, J.C., Benito, R.M.: Measuring political polarization: Twitter shows the two sides of venezuela. Chaos: An Interdisciplinary Journal of Nonlinear Science 25(3), 033114 (2015)
16. Newman, M.E.: Modularity and community structure in networks. Proceedings of the National Academy of Sciences 103(23), 8577–8582 (2006)
17. Nowak, A., Vallacher, R.R., Bui-Wrzosinska, L., Coleman, P.T.: Attracted to conflict: A dynamical perspective on malignant social relations. Understanding social change: Political psychology in Poland pp. 33–49 (2006)
18. Peixoto, T.P.: Hierarchical block structures and high-resolution model selection in large networks. Physical Review X 4(1), 011047 (2014)
19. Rosenthal, S., Farra, N., Nakov, P.: Semeval-2017 task 4: Sentiment analysis in twitter. In: Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017). pp. 502–518. Association for Computational Linguistics, Vancouver, Canada (August 2017), `http://www.aclweb.org/anthology/S17-2088`
20. Stewart, I., Arendt, D., Bell, E., Volkova, S.: Measuring, predicting and visualizing short-term change in word representation and usage in vkontakte social network. In: Proceedings of The 11th International AAAI Conference on Web and Social Media (ICWSM-17). pp. 33–42 (2017)
21. Volkova, S., Chetviorkin, I., Arendt, D., Durme, B.V.: Contrasting public opinion dynamics and emotional response during crisis. In: SocInfo (2016)