

Political Polarization and Popularity in Online Participatory Media: An Integrated Approach

David Garcia
Chair of Systems Design
ETH Zurich
dgarcia@ethz.ch

Fernando Mendez, Uwe Serdült
ZDA, Universität Zürich
{Fernando.Mendez,
Uwe.Serdult}@zda.uzh.ch

Frank Schweitzer
Chair of Systems Design
ETH Zurich
f Schweitzer@ethz.ch

ABSTRACT

We present our approach to online popularity and its applications to political science, aiming at the creation of agent-based models that reproduce patterns of popularity in participatory media. We illustrate our approach analyzing a dataset from *Youtube*, composed of the view statistics and comments for the videos of the U.S. presidential campaigns of 2008 and 2012. Using sentiment analysis, we quantify the collective emotions expressed by the viewers, finding that democrat campaigns elicited more positive collective emotions than republican campaigns. Techniques from computational social science allow us to measure virality of the videos of each campaign, to find that democrat videos are shared faster but republican ones are remembered longer inside the community. Last we present our work in progress in voting advice applications, and our results analyzing the data from *choose4greece.com*. We show how we assess the policy differences between parties and their voters, and how voting advice applications can be extended to test our agent-based models.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and behavioral sciences—*Psychology, Sociology*; H.1.2 [Information Systems]: Models and principles —*User/machine Systems*

Keywords

Popularity, politics, agent-based modeling, emotion

1. INTRODUCTION

Online participatory media, such as social networking sites or discussion forums, play a key role in current political campaigns, and offer the chance to gather large datasets of voter expression and behavior. For example, the messages posted by thousands of *Twitter* users allowed the analysis of the differences in social interaction between parties in the U.S. elections [4, 5]. A common feature of traditional mass media and online participatory media is that both offer politicians a way to reach a very large numbers of users in very

short time. But the main difference is that in the latter, users can also reach large numbers of other users by posting publicly accessible comments. Commonly, we see collective patterns that are emergent phenomena resulting from the interaction of many users. One of these is the popularity of online content, which is difficult to explain by analyzing the “representative user”, due to the inherent heterogeneity of the users of participatory media. Despite these individual differences, the dynamics of votes in *Digg* [1], and of video views in *Youtube* [6, 25], show the existence of statistical regularities of content popularity.

A common approach is to measure online popularity through the amount of users that viewed some content. Because user interaction plays a fundamental role in online communities, the sheer amount of views is not able to grasp other relevant features of popularity. Some examples are how fast the content is shared among users, or the collective emotions expressed in open discussions. In particular, user emotional expression and positive/negative votes, such as *likes* and *dislikes*, are essential for the political sciences. In addition to the number of views or votes, this information contains new levels of complexity. For example, the patterns of positive and negative votes for the content posted in *Reddit*, a social bookmarking website, reveal that the way users vote depends on the votes given by other users [31]. This kind of emotional influence goes beyond objective information sharing, and require psychological explanations.

In our approach, we aim at a unified view of online popularity that links collective and individual behavior, applying it to political parties and candidates. Our statistical analysis of user behavior in online communities reveals robust patterns, such as the diffusion of popular content and the collective emotions expressed towards it. One step further, we also aim at reproducing these patterns by agent-based models based on emotional influence [23]. Agent-based models are starting to be used as a tool in social psychology and emotion research [24]. We have applied these kind of models to reproduce certain patterns of emotional interaction in chatrooms [9], and product reviews [10]. Agent-based models provide a tractable link between the collective human behavior and the interactions of individuals, which is of particular value for political sciences [18]. Agent-based models are based on assumptions that describe the individual behavior and the interaction between agents. These assumptions need to be empirically testable in order to be integrated in a larger scientific perspective, and to be applied in real-world applications. Having access to the online traces of users in social networks is usually not enough to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PLEAD’12, November 2, 2012, Maui, Hawaii, USA.

Copyright 2012 ACM 978-1-4503-1713-9/12/11 ...\$15.00.

design the model at a fine-grained level. In our approach we go beyond the publicly visible layer of the user, integrating other data sources of individual behavior. With respect to political applications, we plan to test our models with data on individual activity in voting advice applications (VAAs). VAAs match the policy preferences of real voters (who visit the website) with that of candidates and/or political parties (that are already encoded in the online system), through the voter’s answers to a set of predefined questions. While designed to provide advice to the users, VAAs also provide an experimental platform to gather data on voter preferences, which cannot be easily retrieved through online traces or election surveys.

In this article, we present our approach to popularity of political campaigns, and how we can apply our findings in political sciences. We illustrate our work in progress by quantifying the success of the campaigns in **Youtube** for the U.S. presidential elections of 2008 and 2012. Our results show how we quantify the content diffusion of a campaign and the collective emotions towards candidates. We follow with a description of our integrated approach, composed of data analysis from online participatory media, agent-based modelling of social interaction in online communities, model validation, and applications to VAAs. The article closes presenting how we can test the patterns of user interaction in political contexts through VAAs, and how such applications can be improved by our models and statistical analysis.

2. YOUTUBE CAMPAIGNS

2.1 Data collection and emotion classification

We want to illustrate our work in progress by showing a relevant example of the application of our approach to the online popularity of political campaigns, in particular to the **Youtube** channels of the candidates for the US presidential elections in 2008 and 2012. For this application, we have retrieved information for all videos in the official channels for Barack Obama (`barackobamadotcom`), Mitt Romney (`mittromney`), and John McCain (`johnmccaindotcom`). In total, 2865 videos were available at the date of the data retrieval, June 4th 2012. For each video, the following data was available: i) time series of views since its creation (more than 190 million views), ii) amount of likes and dislikes (1024402/249986), and iii) comments with text (244010). We divide our dataset in four subsets referring to each of the four studied campaigns. The Obama2008 and McCain2008 subsets contain all the videos in the `barackobamadotcom` and `johnmccaindotcom` channels up to the 2008 presidential elections day. The Obama2012 and Romney2012 subsets contain all the videos in the `barackobamadotcom` and `mittromney` channels created up to a year before the data retrieval date.

The programming interface of **Youtube** limits the access to up to 1000 comments per video, but provides the total number of comments for each video. From this we infer that 70 of the videos of our dataset (2.5% of the total), have more than 1000 comments. The text of each comment has been processed using SentiStrength [28], a sentiment analysis tool for the extraction of emotional content from short texts. SentiStrength is the state-of-the-art tool for lexicon-based analysis of short, informal text. This is the case of **Youtube** comments, for which SentiStrength’s accuracy is above 88% when classifying positive from negative text [27].

In general, it achieves human accuracy, i.e. it can estimate emotional content from text with a similar error rate as a human reader. It has been recently used for the analysis of emotional expression in **Yahoo answers** [14], **Twitter** trends [26], and chatroom communication [9]. SentiStrength uses a lexicon of emotional-bearing terms combined with the detection of negations, amplifiers and diminishers, providing two values, a positive score $p \in [+1, +5]$, and a negative score $n \in [-1, -5]$. As the text of **Youtube** comments is very short (max. 500 characters), we aggregate both values to a measure of polarity. A comment c is classified as positive ($e_c = +1$) if $p + n > 0$, negative ($e_c = -1$) if $p + n < 0$, or neutral ($e_c = 0$) if $p = n$ and both have an absolute value lower than 4. This way, comments with high and equal positive and negative scores are difficult to classify, as we assume that the text is too short to express the coexistence of emotional states. In fact, only 612 comments were detected as the pathological cases of $[+4, -4]$ or $[+5, -5]$, which we discard from our analysis as they form less than 0.3% of the total.

2.2 Collective emotions in political discussions

In addition to the emotional expression in user comments, we also study collective emotions through the likes and dislikes for the videos in our **Youtube** dataset. Each video has a fixed amount of likes and dislikes at the moment of the data retrieval. The distribution of these values per campaign provide information about how positive or negative is the collective response of the **Youtube** community. There is no limit in the amount of likes and dislikes a video can receive, so their values are not bounded from the right. Furthermore, these distributions show a large variance, ranging from 1 to 10 million. For this reason we choose to analyze these distributions in a logarithmic scale, as shown in Fig. 1. These plots already allow us to discard that they follow power-law or log-normal distributions, requiring more data to infer some possible underlying parametric distribution.

Fig. 1 helps us to detect differences in the way users responded to the set of videos of a political campaign. The Obama2008 campaign had a clearly larger average amount of likes than amount of dislikes, while these two values are comparable for the Obama2012 campaign. As the 2008 videos have been available for about four years, the larger amount of likes might have its source in later likes coming after Obama’s victory. Our analysis will continue after the presidential elections of 2012 to assess if this pattern of positivity emerges for the winning candidate, or if it might be taken as support for the detection of different collective emotions.

We aggregate the emotions expressed in the set of comments of a video C_v by calculating the ratios of positive $P_v = \frac{\sum_{c \in C_v} e_c = +1}{|C_v|}$, negative $N_v = \frac{\sum_{c \in C_v} e_c = -1}{|C_v|}$, and neutral $U_v = \frac{\sum_{c \in C_v} e_c = 0}{|C_v|}$ comments. These three values allow us to map each video in a space that can equally represent videos that did not elicit collective emotional responses (high U_v), created positive or negative collective emotions (high P_v or high N_v), or created polarized emotional reactions (both high P_v and N_v). In Fig. 2 we plot each video as a point inside a triangle with a distance to the vertices inversely proportional to U_v for the upper vertex, N_v for the lower left vertex, and P_v for the lower right vertex. Given the overall ratios of emotions in the retrieved comments $\bar{P}, \bar{N}, \bar{U}$, we can perform nonparametric statistical tests to determine

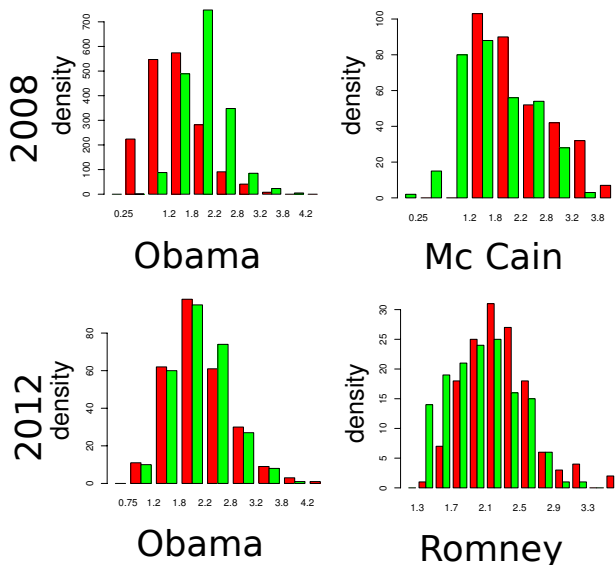


Figure 1: Distributions of the logarithm of the amount of likes (green) and dislikes (red) for the videos in the four studied campaigns.

whether the comments of a video reflect a collective emotional state in the community. This task consist in three χ^2 tests at the 95% confidence interval:

1. Testing U_v versus \bar{U} : if we cannot reject the hypothesis that $U_v \simeq \bar{U}$, the video did not elicit an emotional response (black). If we conclude $U_v > \bar{U}$, the video elicited an underemotional response (gray), for example an exchange of links to other videos or websites. If we conclude $U_v < \bar{U}$, the video created a collective emotional response, and we proceed with the next two tests.
2. Testing P_v versus \bar{P} : if we find $P_v > \bar{P}$, we conclude that the video produced a positive collective emotion.
3. Testing N_v versus \bar{N} : if we find $N_v > \bar{N}$, we conclude that the video produced a negative collective emotion.

If a video passes the first and the second test, but not the third, the video created only a positive collective emotion (green). If it passes the first and the third, but not the second, it created only a negative collective emotion (red). If it passes the first test but not the second nor the third, we know that the video created a collective emotional response, but not of which type it is (yellow). Finally, if a video passes all three tests, we conclude that the video elicited a bipolar collective emotion, which we do not find for this dataset.

Focusing on the 2008 campaigns, the emotions expressed in the comments of Barack Obama's videos were clearly skewed towards positive expression, having numerous instances of positive collective responses and no case of a negative collective emotion. On the other hand, John Mc Cain's videos were triggering much more negative discussions in the few videos uploaded in his channel, having some cases of negative collective emotions. Comparing the two Obama campaigns we can see that the overall pattern of emotional expression in 2012 starts to be similar to the one seen in 2008, but this time two videos elicited collective negative emotions. In the case of Mitt Romney's channel, we

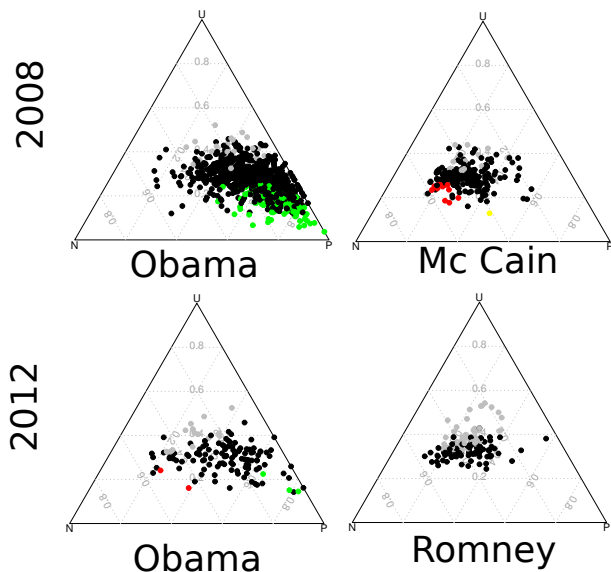


Figure 2: Triangular representation of the collective emotions in the comments for the videos in the four studied campaigns.

did not find any case of collective emotions, but the distribution of emotional expression seems similar to John Mc Cain's without reaching high values of negativity. This kind of quantification of the type of emotional response of the Youtube videos of a political campaign opens the question of whether it is always desirable for a politician to attract large attention to these videos. Popularity does not need to be positive popularity, and online campaigns can indeed encourage negative discussions about the candidate.

2.3 Measuring the 'virality' of political campaigns

The emotions expressed in the comments of a video gave us insights about the collective emotions elicited by the video, but the available data only allowed us to look into the stationary properties of such expression, such as time aggregates. As we mentioned before, popularity is not only composed of stationary metrics like the amount of views, but dynamical properties of the activity of the community. For example the 'virality' of a political campaign can be defined as an analogy to viral marketing, measuring the diffusion of information about the campaign and the adoption by other users in the online community [17]. The online success of a political campaign lies on this information exchange between users. In this section we explore the time series of the amount of views for the videos of the Obama2012 and Romney2012 campaigns, looking for dynamical traces of the success of their campaigns. We restrict our analysis to the Obama2012 and Romney2012 campaigns, focusing on the videos uploaded after February 25th 2012, 100 days before the retrieval date. The reason for this is that the historical data available in the Youtube page of the videos includes time series with a maximum resolution of 100 points, and we would like to investigate the activity of users in the fastest time scale possible, i.e. days in this case. This left us 1310 videos for the Obama2012 campaign and 138 videos for the Romney2012 campaign.

We define $N_v(t)$ as the the total amount of views for a video at time t . For each day, $n_v(t)$ is the amount of views

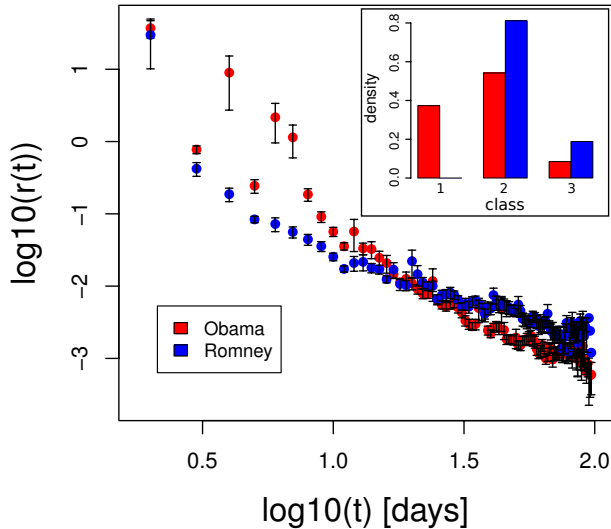


Figure 3: Average growth rates of the videos of the Obama2012 and Romney2012 campaigns. Error bars show standard error over the available videos for each campaign. Inset: Distribution of the classes of view response functions.

that happened only during the day t after the upload of a video, so $N_v(t) = \sum_{t'=1}^t n_v(t')$. We apply a statistical model for the analysis of the complex dynamics of **Youtube** users [6], classifying the different responses of the community through the time series $n_v(t)$. This technique explores two different facets of the collective response to the videos. The first one is the nature of the initiation of the response, which can be *exogenous* if there was some central mechanism that promoted the video from its creation, for example being featured by **Youtube**. An *endogenous* response would be produced by the interaction of a large amount of users, bringing up videos previously unknown to the larger community. The second feature we explore is the criticality of the response, which we use as a way to study the 'virality' of the video. A video elicits a *critical* response if **Youtube** users share and recommend the video at a higher rate than they forget about it, leading to a long response in the views. On the other hand, this response is *subcritical* when the sharing rate is lower than the decay of attention to the video, leading to a very rapid decreasing amount of views. The statistical model introduced in [6] compares the total amount of views of a video, N_v , with the maximum amount of views in a single day, n_v^{max} . This analysis classifies the type of collective response to each video into three classes: exogenous subcritical (1), exogenous critical (2), and endogenous (3).

The inset of Fig. 3 shows the rescaled distributions of the different response classes for the videos of the Obama2012 and Romney 2012 campaigns. In both campaigns we notice that more than 80% of the videos are in the exogenous classes 1 and 2. This is not surprising, as the channels of the presidential candidates should receive a special treatment in **Youtube** due to their public interest, and the videos uploaded in them are very likely to be featured or appear in the main pages of **Youtube**. This poses a very different scenario as in the case of user uploaded videos, where the vast majority cannot receive an external influence that triggers an exogenous response.

Focusing on the Romney2012 campaign, we notice that we did not find any video of class 1 (exogenous subcritical)

within the more than 100 videos uploaded in the channel. This means that every time a video was uploaded and featured in the website, the sharing rate of the users involved kept the viewing activity alive in the following days. These videos keep alive in the community for a long time, which is in line with the stronger social interaction of right leaning users in **Twitter** [4]. The case of the Obama2012 campaign is different, as there is a much larger ratio of videos of class 1, meaning that some videos of this campaign did not trigger the sharing response of the users, who forgot about them quickly. Nevertheless, the unscaled counts show that the Obama2012 contained 710 videos of class 2, while the Romney2012 campaign only 112. This teaches us that, in order to get a large amount of critical responses, or 'viral' videos, a campaign needs to produce a large amount of videos, of which some might not reach this point of criticality and have a lower impact in the community as a whole.

The previous analysis allowed us to detect when the videos of a campaign are above or below the threshold of 'virality', but we can still ask the question of whether we can measure the overall virality of the campaign. To do this, we calculate the growth rate of the amount of views per day for each one of the videos of the campaign defined as $r_v(t) = n_v(t)/N_v(t-1)$. The growth rate $r_v(t)$ measures the amount of new views of a video per previous view, serving as an estimator of the infection rate of a video in the community, or in other words, the average amount of new views triggered by each time a user viewed the video before. Previous analysis of the time evolution of these growth rates of video views [25] showed patterns in the way the overall **Youtube** community behaves. In our case, we focus on the videos of the two studied political campaigns, to draw particular pictures for the users involved in the viewing of such videos. To analyze each campaign, we average over the whole set of available videos at time t , taking $r(t) = \langle r_v(t) \rangle$. The time evolution of $\langle r(t) \rangle$ contains key information on the degree of sharing induced by each campaign.

Fig. 3 shows the time evolution of the growth rates of the views for the Obama2012 and Romney2012 campaigns. The growth rate of the views of the videos in the Obama2012 campaign is significantly larger than the ones in the Romney2012 campaign for the first two weeks after the video is posted. This difference is especially high in the first week, when the growth rates of Obama2012 videos are usually above 1, i.e. on average, videos receive more new views every day than in all the previous days. In the case of Romney2012, only the second day brings an amount of views larger than the first, and then all the growth rates are below 1. On the other hand, the slope of the growth rates shows a stronger decay for Obama2012 than for Romney2012, as after the second week, Romney's videos attracted more views in comparison with the previous amount. This difference is in line with the fact that we did not find videos of class 1 in Romney2012, as the videos of this campaign seem to cause more moderated responses that last longer in time than the ones in the Obama2012 campaign. Obama's videos grow explosively in the first two weeks, but they do not keep this growth rate as probably most of the community willing to view them has already done so. Our data retrieval technique did not impose any demographic or geographic bias, but the communities addressed by this videos might have different locations, age and gender distributions. The demographics

of Youtube might play a role in the origin of these results, and their influence is open for future research.

3. AN INTEGRATED APPROACH TO ONLINE POPULARITY

In this section we present our approach to popularity in online participatory media, and its applications to political sciences. Our work is divided in three areas: data analysis from online participatory media (OPM), agent-based modelling of the interaction between users (ABM), and applications to improve and analyze voting advice applications (VAA). VAAs are online applications deployed during electoral campaigns which match the policy preferences of visiting users with those of candidates and/or political parties that are already stored in the system. The crucial point about VAAs for our study is that they are becoming increasingly popular -millions of visitors in many cases- and generate large datasets of mass public opinion (see more detailed discussion in Section 4). Our results in each of these areas (OPM, ABM, and VAAs) is aimed at generating an impact in the other two areas. This requires a special multidisciplinary effort, because each of these areas address different scientific domains. In particular our analysis of OPM is relevant for statistics, our ABMs for computer science, and our VAAs for political sciences.

3.1 Emergence of popular content

The central aim of our research is to understand the emergence of popularity and polarization in online participatory media and study their influence in politics. We have presented our preliminary results on Youtube in Section 2, but we do not plan to restrict our analysis to the videos of U.S. political campaigns. Websites like Reddit and Digg produced public datasets on the way users voted for articles, including those of political content across countries. In addition, specialized websites like `politnetz.ch` can provide useful information about the support achieved by politicians participating in online social networks. All these communities include commenting functionalities that allow the users to discuss about the content posted on them. To extend our understanding of user behavior in such participatory media, we plan to perform text analysis on the user comments, as we did with SentiStrength on the comments for Youtube videos. Different tools can be used to extract particular knowledge from different kinds of text, and the choice of such tools heavily depends on the length, formality and language of the texts.

Our target in this area is to find statistical regularities of popularity that can be compared across communities, and help to predict future behavior. These regularities, commonly known as *stylized facts* are macroscopic patterns that can be observed at the collective level, but are difficult to predict from the behavior of individual users. In particular we will focus in the different patterns of user emotional reaction to individual political topics, parties, and candidates, as well as how the behavior of users changes when they can perceive content popularity. If large enough, the data contained in OPM can also be used to detect topics of special relevance, which is of particular importance for the efficiency of voting advice applications (VAA). In Section 4, we explain how this output of *topic selection* can be used to improve VAAs.

3.2 Modeling user interaction

Most likely, the *stylized facts* found in our analysis of OPM will not be possible to be explained as a superposition of the behavior of individual users. Agent-based models provide a link between these macroscopic effects, and the microscopic interaction among individuals. In this article we have presented preliminary results in the analysis of the popularity of political campaigns in Youtube. The agent-based models that reproduce the *stylized facts* of campaign virality and collective emotions of Section 2 are current work in progress, following our previous approaches for chat communication and product reviews [9, 10]. The design of these ABM will allow the execution of simulations that span the possible scenarios of the community. In addition, some models allow the usage of tools from statistical physics that provide analytical results on the equations that rule the behavior of the agents, and these in turn can be formulated as predictions about the behavior of the social system as a whole.

Each ABM is based on a set of assumptions that need to be empirically testable, so the design of our models will be driven in a way such that VAAs can provide data for *model validation* at the individual level. The assumptions of our models will be phrased as hypotheses that should be tested against the individual user data, and VAAs can be modified and extended to provide this support. Agent-based models can also provide different metrics that summarize properties of the social interaction between users of VAAs, including their patterns or communication and their positioning towards political issues and candidates. VAAs can be improved with recommendation modules that make use of *network metrics* inspired in the social interaction present in our ABM. Different models of social dynamics allow the formulation of network metrics that can, for example, measure the influence among users from different points of view. For example, different assumptions of social behavior can be used for spammer detection in Twitter [11], by applying their related network metrics. Furthermore, analytical results of ABMs can be applied to *mechanism design*, providing possible improvements in the decision space of the designers of online communities. Simulations of these models can deliver insights of the impact of different decisions and policies followed by community managers, for example in order to maximize the impact of political campaigns, or to trigger discussions that allow voters to deepen in their own opinions.

3.3 Measuring individual dynamics

For many users, political leaning is a private concern that is not shared publicly. This leads to numerous problems, like self-selection bias, which need to be accounted for when studying collective expression of political opinions [12]. This *silent majority* can compose a large part of the voting community, and data of this kind is rather scarce. In our work, we want to overcome the limitations of analyzing a *vocal minority* that might not be representative of the voter community as a whole. Our aim is to explore deeper than the public behavior of general users, studying the dynamics of individual users and voters. The high degree of privacy of VAAs, in comparison with comments on participatory media, allow users to freely explore their political convictions without the fear of being exposed. When using a VAA, users allow the usage of their anonymized data for research, but

no user account, email or any other kind of personal data is required to use the platform.

We count with previous implementations of VAAs that attracted significant amounts of users for local and national elections in Greece, Cyprus, London, Scotland, Peru, and Brazil. We plan to develop some extensions of their functionality to enhance the data they produce in order to address the hypotheses generated by our ABMs, as *model validation* is an essential block of the modeling cycle. To explore the individual dynamics of emotions and opinions, we plan to implement three new mechanisms in our VAA platform: i) emotional feedback, ii) user discussions, and iii) user to user invitations. The emotional feedback module will consist on a set of inquiries to the users that request them to provide a self-report of emotions in a Likert scale, in particular related to political topics, candidates, and the output of the VAA itself. The user discussion module will consist on a forum that contains discussions for the different political topics of the VAA, on which the users can make completely anonymous comments. This data will allow us to couple the internal emotions of the users with their political leaning, their voting intention, and their public online expression. In the next section, we explain our previous and current work on VAAs, leaving our work related to ABMs and data visualization for future publications.

4. APPLICATIONS TO POLITICAL SCIENCES

4.1 Voting Advice Applications

VAAs are online applications that gather the policy preferences of the visiting user, matching them to the policy statements of candidates and/or political parties. Developed in The Netherlands in the 1990s, and initially paper-based, VAAs have now proliferated across many European countries during election campaigns. Furthermore, participatory media has further increased the dissemination potential of such tools across the wider electorate. To date it appears that VAAs have become more 'institutionalized' in a number of specific national electoral settings: Belgium, Germany, The Netherlands, and Switzerland. What is striking about the deployment of VAAs in some of these pioneer cases is the sheer magnitude of voting 'recommendations' issued, between 20-40 per cent of the electorate in some elections and around 5 million recommendations in the case of German elections in 2005. These countries share important electoral features, namely the fact that they are multi-party systems. Indeed, it has been argued that VAAs are especially well suited to multi-party electoral settings [22]. The same appears to be true for systems where elections are candidate-centred rather than party based [29, 16].

Another reason for the increasing online popularity of VAAs is their relative simplicity. The mechanism is rather straightforward. Prior to an election, candidates/parties provide their policy positions by filling in a questionnaire containing an extensive set of policy statements. In some cases, the policy positions of candidates/parties are coded by political science experts. When the VAA is launched during the election campaign, citizens can then fill in the same policy questionnaire. An algorithm matches the responses of candidates/parties with the ones provided by the citizens using the tool. The core output is an ordered ranking of candidates/parties according to the degree of similarity to users' preferences. In short, the aim of the tool is to allow

citizens to better define their own subjective, political preferences and to match these with the stated (or expert coded) preferences of candidates or political parties. The end result should be a more informed vote choice among the range of parties/candidates competing in the political space.

4.2 Insights on voter behavior

Our VAA implementations provide substantial datasets on the individual voter and party policy questionnaires. Our analysis of these datasets focuses on one question that has occupied a sub-field of political science for many years: how best to estimate the positions of political parties in the ideological space. Typically, two dominant strategies have been employed to estimate the ideological positions of parties: expert surveys (usually sent to a handful of political scientists) or the content analysis of party manifestos. It is much rarer to find the use of mass surveys to estimate political party positions. VAA-generated datasets are ideally suited to this task for three reasons: i) they typically contain a much more extensive set of policy statements than the surveys conducted by national election teams, ii) the datasets they produce are very large, containing millions of users, and iii) most VAAs also contain supplementary questions, such as the socio-demographic profile of users and their vote intention or party affiliation. With such data it is possible to map the positions of political parties in the ideological space, based on the mass opinions of their supporters.

Fig. 4 shows the ideological space of the recent elections in Greece on May 6th 2012. The scatter plot shows the data provided by the users visiting the Greek elections VAA choose4greece.com, and the expert coding of the political parties. This political map consists of a left-right x axis, and a y axis in which the VAA designers incorporated a new dimension that structured the Greek elections of 2012 -pro versus anti positions to the Greek bailout conditions imposed by the European Union and IMF. Squares in Fig. 4 show the mean position of partisan supporters based on their answers to the policy statements that loaded on to each of the dimensions. Circles represent the expert coding of the corresponding parties, connected to the squares aggregating the users willing to vote for the corresponding party. Party supporters occupy a less polarized position than the academically coded positions of the political parties. Whilst it is to be expected that the average partisan supporter is less consistent across an extensive set of policy statements, it is important to note that all the mean positions of the partisan supporters are in the correct quadrants. Such dynamics could be relevant to broader debates about processes of party dealignment or realignment among political scientists. This result is in line with the increasing trend towards party dealignment and growing numbers of floating voters [32, 16], which in turn is an additional argument to use VAAs as a tool for electoral studies.

The lower polarization of the average party supporters versus the expert-coded policy position of the parties suggests that voters prefer parties that take more polarized positions, and not ones that are necessarily closest to them in the political space. This leads us to a concern about VAAs of relevance for computational science, which is how to produce a best 'match' between voters using the application and the parties/candidates. This is especially important since VAAs can be used in some cases by a significant share of the electorate. In short, the choice of algorithm

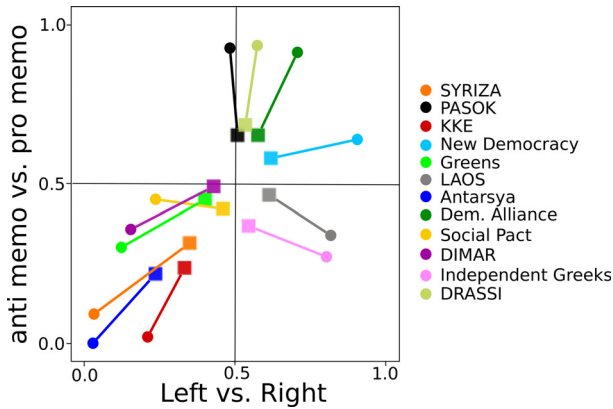


Figure 4: Representation of the two-dimensional projection of the questionnaire responses for parties (circles) and their voters (squares) in the VAA for the first Greek elections 2012. The horizontal axis refers to left vs. right oriented policies, and the vertical axis to the positioning regarding the bailout.

used by VAA designers could have electoral consequences. Presently, all VAAs are based on a social choice theory of democracy [7], a rational model of voter behavior that goes back to Downs’ influential economic theory of democracy [7]. The voter is assumed to compare her stance with that of the candidates/parties on all of the issues and chooses the candidate/party which offers, on balance, the largest number of preferred policies. Over the years these ideas have been formalized into a spatial theory of electoral choice where voters choose the party/candidate whose policy position is most proximate to theirs [33, 30]. These ‘proximity’ models form the theoretical core of a VAA, typically using a City Block measure or, less commonly, a Euclidean metric.

But there are other theories of voting which are predicated on different assumptions. Directional theories conceptualize the policy dimension differently [15, 33, 30], proposing that what matters most is for the voter and candidate/party to be on the correct ‘side’ of the argument. According to directional logic, voters do not tend to distinguish fine policy gradations and are more concerned that their candidate is on the correct side of the argument. Indeed, the theory suggests that, within a boundary of acceptability, voters prefer more extreme candidates. The metric used in a *directional* theory is the scalar product. VAA generated data is ideally suited for testing the predictive power of these competing hypotheses of voting behavior. A preliminary analysis by Mendez [20] finds that directional inspired algorithms perform better in predicting the voter intention of partisan VAA users than proximity metrics in four electoral settings.

4.3 Improving VAAs through social media

A crucial element to VAA design is the selection of policy statements. Policy statements should be chosen that best polarize the parties/candidates as well as the voters using the tools. Here the problem is that the choice of policy statements (and their wording) is not necessarily a neutral affair. Simulation results [32] show that the selection of particular statements has a considerable impact on the voting recommendations that are produced. The configuration of the policy questionnaire could benefit some parties disproportionately. Through their selection of items, it could be argued that VAA designers are helping to reinforce the main-

stream policy agenda. We plan to improve VAAs through a mechanism of *topic selection* based on information retrieval from OPM. While the wording of a statement is a delicate issue that requires political science expertise, the set of topics addressed by the statements can be improved automatically. We plan to implement a feedback system that monitors participatory media, focusing on a large set of possible relevant topics. The system will measure the polarization of user generated content on the topic, and provide a rank of the most significant topics for the VAA. This tool would be a new channel for incorporating policy options that are typically neglected by mainstream media, reshaping users’ perceptions of the political landscape. Such an approach would incorporate elements of the ‘contestatory’ model of democracy, which emphasizes the need for alternative viewpoints to mainstream opinion as a crucial element in a well functioning democracy [21].

With the advent of social network platforms, in particular Facebook and Twitter, VAAs can quickly become a viral affair. Some of this can be gleaned from the thousands of ‘likes’ and ‘tweets’ displayed by VAA sites. Indeed, a VAA can also be considered as a socio-technical system, capable of generating meaningful mass network data. Such data could lead researchers to develop a better understanding of opinion dynamics during online political campaigns. In a first step, we will develop a network component for our VAA that would allow the user (ego) to share obtained results with friends (alteri) and to invite them to participate as well. This approach aims at overcoming the simplified view of a VAA user, which often is limited to its answers to the questionnaire and surveys. The addition of social network and interaction data allows an extended view of the user beyond its relation to party policies.

5. CONCLUSIONS

We have presented our preliminary results on the statistical analysis of popularity in participatory media, through the collective emotions and ‘virality’ of political videos in Youtube. These results shed light on the different mechanisms of emotional interaction and political information diffusion in online communities. We introduced our broader perspective on the topic of popularity in participatory media and its applications to the political sciences, integrating collective and individual dynamics through agent-based models. Our current platforms of VAAs provide novel data that complements what we can extract from the online traces of Internet users. This data can be used to test different hypotheses of voter behavior, and to support the assumptions of our ABMs. Furthermore, future VAAs can include a series of socially-aware tools, like topic selection based on user discussions.

Future analysis of various online communities as well as new instances of VAAs can provide large amounts of high quality data on voter behavior. We expect processes of social comparison, cognitive dissonance and homophily to be at work not only in the real but also in the virtual world [8, 19]. VAA network data can then be analyzed with the help of concepts and indicators well known in social network analysis [3]. Reciprocity within social networks of VAA users would for example be expected to go hand in hand with similar political beliefs and values [2, 13].

6. ACKNOWLEDGEMENTS

Authors acknowledge A. Abisheva for her assistance on data retrieval.

7. REFERENCES

- [1] S. Asur, B. A. Huberman, G. Szabo, and C. Wang. Trends in Social Media: Persistence and Decay. In *ICWSM'11*, Palo Alto, 2011.
- [2] P. F. Berelson, Bernard R. Lazarsfeld and W. N. McPhee. *Voting: A Study of Opinion Formation in a Presidential Campaign*. University of Chicago Press, 1954.
- [3] S. P. Borgatti, A. Mehra, D. J. Brass, and G. Labianca. Network analysis in the social sciences. *Science*, 323(5916):892–895, 2009.
- [4] M. D. Conover, B. Gonc, A. Flammini, and F. Menczer. Partisan Asymmetries in Online Political Activity. *EPJ DataScience*, 1(6):1–17, 2012.
- [5] M. D. Conover, B. Goncalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the Political Alignment of Twitter Users. In *SocialCom'11*, pages 192–199, 2011.
- [6] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences of the United States of America*, 105(41):15649–15653, Oct. 2008.
- [7] A. Downs. An economic theory of political action in a democracy. *The Journal of Political Economy*, 65(2):135–150, 1957.
- [8] L. Festinger. A theory of social comparison processes. *Human Relations*, 7:117–140, 1954.
- [9] A. Garas, D. Garcia, M. Skowron, and F. Schweitzer. Emotional persistence in online chatting communities. *Scientific Reports*, 2:402, May 2012.
- [10] D. Garcia and F. Schweitzer. Emotions in Product Reviews—Empirics and Models. In *SocialCom'11*, pages 483–488, 2011.
- [11] D. Gayo-Avello. Nepotistic Relationships in Twitter and their Impact on Rank Prestige Algorithms. In *Congreso Español de Recuperación de Información*, page 31, Madrid, Spain, 2010.
- [12] D. Gayo-avello. I Wanted to Predict Elections with Twitter and all I got was this Lousy Paper - A Balanced Survey on Election Prediction using Twitter Data. arXiv:1204.6441v1, 2012.
- [13] R. Huckfeldt. Interdependence, density dependence, and networks in politics. *American Politics Research*, 37(5):921–950, 2009.
- [14] O. Kucuktunc, B. B. Cambazoglu, I. Weber, and H. Ferhatosmanoglu. A large-scale sentiment analysis for Yahoo! answers. In *WSDM'12*, page 633, 2012.
- [15] D. Lacy and P. Paolino. Testing proximity versus directional voting using experiments. *Electoral Studies*, 29:460–471, 2010.
- [16] A. Ladner and J. Pianzola. Do voting advice applications have an effect on electoral participation and voter turnout? *Electronic Participation*, pages 211–224, 2010.
- [17] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web*, 1(1):39, May 2007.
- [18] J. Lorenz. Zur Methode der agenten-basierten Simulation in der Politikwissenschaft am Beispiel von Meinungsdynamik und Parteienwettstreit. In *Jahrbuch für Handlungs- und Entscheidungstheorie. Band 7: Experiment und Simulation*. VS Verlag für Sozialwissenschaften, 2012.
- [19] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:pp. 415–444, 2001.
- [20] F. Mendez. Matching voters with political parties and candidates: An empirical test of four algorithms. *International Journal of Electronic Governance*, (Special issue on Voting Advice Applications and the State of the Art), forthcoming.
- [21] P. Pettit. Democracy, electoral and contestatory. In I. Shapiro and S. Macedo, editors, *Designing Democratic Institutions*. Nomos, NYU Press, New York, 2000.
- [22] A. Ramonaitė. Voting advice applications in lithuania: Promoting programmatic competition or breeding populism? *Policy and Internet*, 2(1):117–141, 2010.
- [23] F. Schweitzer and D. Garcia. An agent-based model of collective emotions in online communities. *The European Physical Journal B*, 77(4):533–545, 2010.
- [24] E. R. Smith and F. R. Conrey. Agent-based modeling: a new approach for theory building in social psychology. *Personality and social psychology review : an official journal of the Society for Personality and Social Psychology, Inc*, 11(1):87–104, Feb. 2007.
- [25] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Communications of the ACM*, 53(8):80, Aug. 2010.
- [26] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment in Twitter events. *Journal of the American Society for Information Science and Technology*, 62(2):406–418, 2011.
- [27] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment strength detection for the social web. *Journal of the American Society for Information Science and Technology*, 63(1):163–173, Jan. 2012.
- [28] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas. Sentiment strength detection in short informal text. *Journal of the American Society for Information Science and Technology*, 61(12):2544–2558, Dec. 2010.
- [29] J. Thurman and U. Gasser. Three case studies from switzerland: Smartvote. *Berkman Center Research Publication*, No. 2009-03.3, 2009.
- [30] M. Tomz and R. P. Van Houweling. Candidate positioning and voter choice. *American Political Science Review*, 102(03):303–318, 2008.
- [31] P. Van Mieghem. Human Psychology of Common Appraisal: The Reddit Score. *IEEE Transactions on Multimedia*, 13(6):1404–1406, Dec. 2011.
- [32] S. Walgrave, M. Nuytemans, and K. Pepermans. Voting aid applications and the effect of statement selection. *West European Politics*, 32(6):1161–1180, 2009.
- [33] A. Westholm. Distance versus direction: The illusory defeat of the proximity theory of electoral choice. *The American Political Science Review*, 91(4):865–883, 1997.