

RESEARCH

Open Access



Forecasting elections results via the voter model with stubborn nodes

Antoine Vendeville^{1,2,3*} , Benjamin Guedj^{1,2,3,4} and Shi Zhou^{1,2,3}

*Correspondence:
a.vendeville@ucl.ac.uk
³ Centre for Artificial
Intelligence, University
College London, London, UK
Full list of author information
is available at the end of the
article

Abstract

In this paper we propose a novel method to forecast the result of elections using only official results of previous ones. It is based on the voter model with stubborn nodes and uses theoretical results developed in a previous work of ours. We look at popular vote shares for the Conservative and Labour parties in the UK and the Republican and Democrat parties in the US. We are able to perform time-evolving estimates of the model parameters and use these to forecast the vote shares for each party in any election. We obtain a mean absolute error of 4.74%. As a side product, our parameters estimates provide meaningful insight on the political landscape, informing us on the proportion of voters that are strong supporters of each of the considered parties.

Keywords: Elections, Voter model, Opinion dynamics, Markov chains, Social networks

Introduction

For decades, modern democratic societies have been polling populations to try and track the popularity of elections candidates and members of governments. Those are often conducted by means of phone, online or even in person surveys, which can be very time-consuming and usually suffer from limited sample sizes and bias—*e.g.* respondents with controversial views might be reluctant of sharing them. This is why different methods are being investigated nowadays. With the rapid growth of online social platforms such as Facebook or Twitter, any individual can now publicly express their views and opinions, adding to an evergrowing pool of directly accessible data. This has open the door for a new avenue of research, that seeks to use this precious resource to forecast polls and election results without having to survey the population.

As of today, most efforts have focused on applying machine learning methods such as sentiment analysis to evaluate public opinion through samples of Twitter data and try to predict the outcome of democratic processes around the globe (Saleiro et al. 2016; Garcia et al. 2018; Grimaldi et al. 2020). The quality of predictions spans a rather wide range and numerous voices have expressed concerns over these methods, arguing that there are multiple factors at play that may alter their reliability (Gayo-Avello 2012; Jungherr et al. 2017). This is why in this work we propose a novel method that does not rely on data analysis but rather uses the authentic and official results of previous elections to perform estimation for future ones.

More precisely, we consider the well-known voter model for opinion dynamics. A population of connected nodes form a graph where some of them are in state 0 and some others in state 1. Nodes can then randomly change state over time following the distribution of others' states. Nodes are usually meant to represent users on a social network and states their opinions or views. This model then allows to describe in a simple and intuitive manner social dynamics where people are divided between two parties and form their opinion by observing that of others around them. A previous work of ours was dedicated to the theoretical study of this model in the specific case where everyone is influenced by everyone else and some users are stubborn and never change opinion (Vendeville et al. 2020). Notably, we provided closed-form expressions for the distribution of opinions at any point in the process and convergence time to equilibrium.

This paper is a follow-up of that work, as we apply our previous findings to develop a novel method that can be used to forecast the results of any election. We look at both general elections in the United Kingdom and presidential elections in the United States. In each case, we consider the evolution of the share of popular votes for each of the two major parties¹ as a realisation of the voter model and perform time-evolving estimation of optimal parameters. This allows us to obtain a theoretical distribution for the number of seats or votes, from which we draw the expected result of future elections. We compare with real-life outcomes to assess the viability of our approach.

Related literature

A number of research projects have focused on applying machine learning algorithms to Twitter data in order to forecast opinion poll results or election outcomes. We discuss some of them here and refer the interested reader to Gayo-Avello (2013) and Phillips et al. (2017) for more in-depth reviews of the literature. A pioneer work in this area was that of Tumasjan et al. (2011) whose model achieved a mean average error (MAE) of 1.65% when predicting results of the 2009 German federal election. Authors used Twitter mention counts as an direct indicator of a candidate's popularity, a method that has been considered by several other works as well, often in combination with a sentiment analysis of tweets content (O'Connor et al. 2010; Saleiro et al. 2016; Garcia et al. 2018; Grimaldi et al. 2020; Fink et al. 2013; Huberty 2013; Caldarelli et al. 2014; Thapen and Ghanem 2013). In particular, Garcia et al. (2018) achieved 90% accuracy in predicting the top two candidates in various municipalities during Brazilian municipal elections, and Saleiro et al. (2016) achieved a MAE of 0.63% when trying to predict opinion poll results during the Portuguese bailout (2011–2014).

The relevance of such approaches has however been questioned by a number of authors (Fink et al. 2013; Huberty 2013; Caldarelli et al. 2014; Thapen and Ghanem 2013; Jungherr et al. 2012, 2017; Gayo-Avello 2012). Jungherr et al. (2012) showed that merely changing the timeframe of forecast in the work of Tumasjan et al. (2011) would invalidate the results. Fink et al. (2013) found that the use of Twitter mentions mirrored actual popularity of only some of the candidates but not all of them. Jungherr et al. (2017) argued that mentions count, used in most of the works cited above, show evidence of

¹ Conservative and Labour in the UK, Republican and Democrat in the US.

attention to politics rather than support to the actual candidates. This is why researchers often combine mentions count with sentiment analysis algorithms, but even these can have trouble detecting and correctly interpreting all subtleties of the human language. This particular concern, has been raised by several authors (Huberty 2013; Caldarelli et al. 2014; Gayo-Avello 2012). Self-selection, i.e. the fact that people choose whether to express their views online or not, may also bias results. Add to it the rife presence of bots on the Twitter platform which makes it delicate to assess whether the online population is an accurate representation of the real one.

Some researchers have thus considered different avenues, drawing features from the Twitter user graph topology (Dokoochaki et al. 2015), hashtags co-occurrences (Bovet et al. 2018) or even discarding the social platform entirely and using fluctuations of the Pound to forecast the popularity of the Conservative party in the UK (Usher and Dondio 2020). Integrating in this line of works, we build a model that does not rely on Twitter but rather uses official results of previous elections to guess the outcome of future ones. Our model is a variant of the celebrated voter model, where nodes on a graph are in one of two possible states and repeatedly update their beliefs to agree with other nodes chosen at random. It was introduced independently by Holley and Liggett (1975) and Clifford and Sudbury (1973) in the context of particles interaction. They proved that consensus is reached, i.e. that every node is eventually in the same state, on the infinite \mathbb{Z}^d lattice. Several works have since looked at different network topologies: complete graphs (Hassin and Peleg 2002; Sood et al. 2008; Perron et al. 2009; Yildiz et al. 2010), Erdős-Rényi random graphs (Sood et al. 2008; Yildiz et al. 2010), scale-free random graphs (Sood et al. 2008; Fernley and Ortgieste 2019), and other various structures (Yildiz et al. 2010; Sood et al. 2008). Variants where nodes deterministically update to the most common state amongst their neighbours have also been studied (Chen and Redner 2005; Mossel et al. 2014).

In this paper we consider the specific case where stubborn nodes who never switch state are present in the graph. Such nodes may for example represent lobbyists, politicians or activists, i.e. entities looking to lead rather than follow and who will not easily change side. One of those placed within the network can singlehandedly change the outcome of the process (Mobilia 2003; Sood et al. 2008). If several of them are present on both sides, consensus is usually not reachable and instead the distribution of states converges to an equilibrium in which it fluctuate indefinitely (Mobilia et al. 2007; Yildiz et al. 2013). Recently, Mukhopadhyay et al. (2020) considered nodes with different degrees of stubbornness and show that time to reach consensus grows linearly with their number. Klamser et al. (2017) studied the effect of stubborn nodes on a dynamically evolving graph, and show that the two main factors shaping their influence are their degrees and the dynamical rewiring probabilities. Finally, in our previous work we developed closed-form formulas for the distribution of opinion at any step and convergence time to equilibrium in the case where stubborn nodes are present in a strongly connected network (Vendeville et al. 2020).

Our contributions In this paper we propose a new model for the forecast of elections outcome, based on official results of previous elections. Our method is based on the voter model with stubborn nodes and uses theoretical results developed in a previous work of ours (Vendeville et al. 2020). We apply it to the United Kingdom general elections and in

the United States presidential elections and achieve an MAE of 4.74%. To the best of our knowledge this is the first time such work is conducted. All code used is available online.

Theoretical background

Here we present the mathematical framework behind our forecasting method. In the traditional voter model, we consider a group of n nodes labelled $1, \dots, n$ who are each in state 0 or 1. These states are prone to change over time and we let $x_i(t)$ denote the state of node i at time t . Each node has access to the state of some of the others, called its neighbours. Nodes can then be seen as forming a graph of size n , with an edge from j to i if and only if i has access to the state of j . Here we consider this graph to be a clique with unweighted edges and no self-loops. Thus each node accounts for the state of every other, except their own, with no particular preference. The process then unfolds as follows. Starting with a given initial distribution of states, an independent exponential clock of parameter 1 is associated to each node. Whenever a clock rings, the concerned node changes its state to that of one of its neighbours selected uniformly at random—or equivalently, chooses its new state by sampling the distribution of its neighbours' states.

We let $N_1(t)$ denote the number of state-1 holders at time t ; it will be our quantity of interest. Note that the number of state-0 nodes at time t is given by $n - N_1(t)$. We assume $N_1(0)$ is fixed and let n_1 denote its value. We are interested in the particular situation where some of the nodes are stubborn, that is never change state, and we describe the evolution of $N_1(t)$ over time. We denote by s_0 and s_1 the numbers of stubborn state-0 and state-1 nodes respectively and require at least one of them to be strictly positive. To this end we define

$$S_n = \{(a, b) \in \{0, \dots, n\}^2 : 0 < a + b \leq n\} \tag{1}$$

and require $(s_0, s_1) \in S_n$. We write $[m_{ij}]_{i,j}$ to denote the matrix with entry m_{ij} in the i -th row and j -th column and let e^M denote the exponential of any matrix M .

Because s_0 and s_1 nodes will always be in respective states 0 and 1 no matter what, $N_1(t)$ is comprised between s_1 and $n - s_0$ for all t . The idea behind our analysis is that it describes a birth-and-death process over the state-space $\{s_1, \dots, n - s_0\}$ with transition rates, for all $s_1 \leq k \leq n - s_0$,

$$\begin{cases} q_{k,k-1} = (k - s_1)(n - k)/(n - 1) \\ q_{k,k+1} = k(n - k - s_0)/(n - 1) \\ q_{k,k} = -q_{k,k-1} - q_{k,k+1}. \end{cases} \tag{2}$$

Indeed to move from state k to $k - 1$ we need a non stubborn state-1 node to adopt the state of an state-0 node. There are $k - s_1$ non stubborn state-1 nodes and for each of these, a proportion $(n - k)/(n - 1)$ of the others is in state 0, hence $q_{k,k-1} = (k - s_1)(n - k)/(n - 1)$. We obtain $q_{k,k+1}$ via an analogous reasoning and define $q_{k,k} = -q_{k,k+1} - q_{k,k-1}$. Since the process only evolves by unit increments or decrements, $q_{k,j} = 0$ if $j \notin \{k - 1, k, k + 1\}$. As expected we have $q_{s_1, s_1 - 1} = 0$ and $q_{n - s_0, n - s_0 + 1} = 0$. Finally we let $Q = [q_{ij}]_{i,j}$ denote the transition rate matrix and e^{tQ} the exponential of tQ defined by $e^{tQ} = \sum_{r=0}^{\infty} (tQ)^r / r!$ for any $t > 0$. From there we are able to compute the distribution of $N_1(t)$ and its expected value at any point in time.

Theorem 1 *Let Q be the matrix with entries described in (2) and let $N_1(0) = n_1$ be given. Assuming $(s_0, s_1) \in S_n$ is the repartition of stubborn nodes, the probability for N_1 to equal k at time t is*

$$p_{n_1,k}(t) := [e^{tQ}]_{n_1,k}. \tag{3}$$

Hence,

$$\mathbb{E}N_1(t) = \sum_{k=s_1}^{n-s_0} k p_{n_1,k}(t). \tag{4}$$

is the expected number of state-1 nodes at time t .

Because there are stubborn agents in both camps, consensus is never reached and instead, the system indefinitely fluctuates within a state of equilibrium. More precisely, $N_1(t)$ converges in distribution as $t \rightarrow \infty$. The limiting distribution is called *stationary* and we denote it by $\pi = (\pi_{s_1}, \dots, \pi_{n-s_0})$. We would like to know if the political system we consider can be considered to be within such state. To this end, the long term expectation of $N_1(t)$ is given by the following theorem.

Theorem 2 *Assuming $(s_0, s_1) \in S_n$ is the repartition of stubborn agents, $N_1(t)$ has a unique stationary distribution $\pi = (\pi_{s_1}, \dots, \pi_{n-s_0})$ and thus the expected number of opinion-1 holders converges to*

$$\mathbb{E}\pi = n \frac{s_1}{s_0 + s_1}. \tag{5}$$

The theory has been developed in our previous work (Vendeville et al. 2020) to which we refer the interested reader for more details and proof of Theorems 1, 2. We also provide a closed-form formula for the computation of convergence time.

Setup

We use the official database of the United Kingdom general elections results from 1922 onwards, published by the House of Commons (Audickas et al. 2020), as well as results for presidential elections in the United States from 1912 onwards, manually collected from Wikipedia.² Each time we are interested in the percentage of popular votes won by each of the two major parties—Conservative and Labour in the UK, Republicans and Democrats in the US. We assume these quantities correspond to pointwise observations of independent realisations of the voter model. The result of each election can then be forecast via Theorem 1, provided we have an estimate of the quantity of stubborn nodes (s_0, s_1) . Thus, our analysis is done in two steps: first we make for each election an estimate of (s_0, s_1) based on previous results, then Eq. 4 gives us the expected value for the coming election that we use as a predictor.

For the sake of clarity we present our method in the UK case, but note that it directly translates to the US case. Different parties are present, the two major ones being

² https://en.wikipedia.org/wiki/United_States_presidential_election#Popular_vote_results.

Conservative³ and Labour, the rest including Liberal Democrats or Social Nationalists amongst others. Because our model applies to a two-sided situation only, we cannot consider all of them at once. Thus, we aggregate all non-Conservative parties under the label 0 while Conservatives are attributed label 1. We let x_i denote the percentage of seats won by Conservatives on the i th election, rounded to the nearest integer because our model cannot account for decimal values. In addition we let t_i denote the elapsed time, in years, since the starting point 1922. There have been $m = 27$ elections total, with the last one taking place in 2019. Thus $t_1 = 0$ and $t_m = 2019 - 1922 = 97$. We let x_m denote the percentage of seats won by the conservatives in 2019. To concur with our theoretical framework we consider one seat won by the Conservatives (resp. non-Conservatives) as the observation of an node being in state 1 (resp. 0) amongst $n = 100$ of them. The x_i 's then correspond to pointwise observations at times t_i 's of a realisation of the process $N_1(t)$ described in Sect. 3. All the reasoning described here and in the following will also be applied independently in the cases Labour versus non-Labour, Republican versus non-Republican (US) and Democrat versus non-Democrat (US).

Methodology

To be able to use Theorem 1 to make predictions, we first need to estimate the proportion of potential stubborn nodes in the population, that is the percentage of votes which are guaranteed either for Conservatives or other parties. Let s_0 denote the number of stubborn state-0 (non-Conservative) nodes and s_1 that of state-1 (Conservative) ones. We look for the values (s_0^*, s_1^*) that maximise the log-likelihood of the observed data (x_1, \dots, x_m) under the assumption that those were generated via a realisation of the voter model. Let's say we want to predict results for the i th election. Because we need at least two datapoints to make an estimation, we require $3 \leq i \leq m + 1$. Following the notations introduced in Sect. 3 we let $p_{k,l}^{(s_0,s_1)}(t)$ denote the theoretical probability for $N_1(t)$ to go from k to l in t units of time when there are respectively s_0 and s_1 state-0 and state-1 stubborn nodes. We seek to solve

$$\arg \max_{s_0, s_1} \sum_{j=1}^{i-2} \log \left(p_{x_j, x_{j+1}}^{(s_0, s_1)}(t_{j+1} - t_j) \right). \tag{6}$$

Indeed, $p_{x_j, x_{j+1}}^{(s_0, s_1)}(t_{j+1} - t_j)$ is by definition the probability for Conservatives to win x_{j+1} percent of the votes in the $(j + 1)$ th election knowing they won x_j percent in the j th one. Thus we seek to simultaneously maximise the likelihood of all past elections results. Let $Q^{(s_0, s_1)}$ be the matrix with entries calculated via (2). By Theorem 1, we have that (6) is equivalent to

$$\arg \max_{s_0, s_1} \sum_{j=1}^{i-2} \log \left[e^{(t_{j+1} - t_j) Q^{(s_0, s_1)}} \right]_{x_j, x_{j+1}} \tag{7}$$

The computation of matrix exponential is typically done in cubic time and quickly becomes intractable as the size of the matrix increases. Here however, because we have

³ The dataset also includes in Conservative results: National, National Liberal and National Labour candidates for 1931–1935; National and National Liberal candidates for 1945; National Liberal candidates from 1945 to 1970.

Table 1 Evolution of the estimates for the proportion of stubborn agents (s_0^*, s_1^*) over time

Year	Conservative	Labour	Year	Republicans	Democrats
1924	(62, 38)	(65, 30)	1920	(23, 23)	(44, 42)
1929	(20, 21)	(61, 30)	1924	(15, 21)	(18, 12)
1931	(28, 20)	(55, 30)	1928	(18, 23)	(15, 8)
1935	(1, 5)	(53, 27)	1932	(18, 23)	(18, 11)
1945	(9, 10)	(48, 26)	1936	(16, 18)	(10, 8)
1950	(11, 10)	(26, 17)	1940	(13, 13)	(7, 7)
1951	(13, 12)	(23, 16)	1944	(14, 14)	(9, 8)
1955	(13, 12)	(23, 16)	1948	(15, 15)	(9, 8)
1959	(15, 14)	(22, 16)	1952	(17, 16)	(10, 9)
1964	(16, 15)	(25, 18)	1956	(16, 16)	(11, 10)
1966	(18, 16)	(25, 18)	1960	(16, 16)	(11, 10)
1970	(18, 16)	(24, 18)	1964	(17, 17)	(12, 11)
1974	(19, 17)	(26, 19)	1968	(17, 16)	(10, 10)
1974	(19, 16)	(26, 19)	1972	(17, 16)	(12, 11)
1979	(19, 16)	(26, 19)	1976	(15, 15)	(11, 10)
1983	(20, 17)	(28, 20)	1980	(16, 16)	(12, 11)
1987	(20, 17)	(22, 15)	1984	(16, 16)	(13, 11)
1992	(22, 18)	(21, 14)	1988	(16, 16)	(13, 11)
1997	(22, 18)	(23, 15)	1992	(16, 16)	(14, 12)
2001	(19, 15)	(24, 16)	1996	(15, 15)	(14, 12)
2005	(18, 14)	(24, 16)	2000	(16, 15)	(15, 13)
2010	(17, 13)	(24, 16)	2004	(16, 15)	(15, 13)
2015	(18, 13)	(22, 14)	2008	(16, 16)	(15, 13)
2017	(18, 13)	(22, 14)	2012	(17, 16)	(16, 14)
2019	(19, 14)	(22, 14)	2016	(17, 16)	(16, 14)
2024	(19, 14)	(24, 15)	2020	(18, 17)	(16, 14)

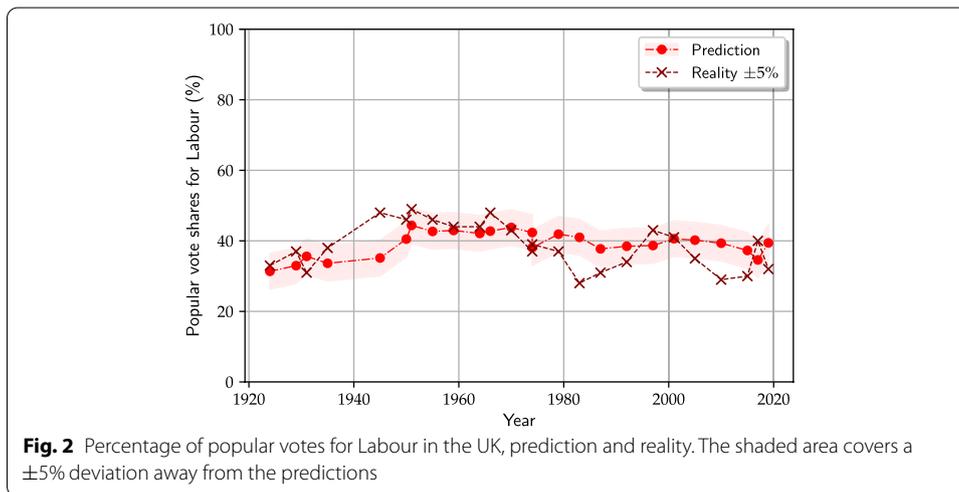
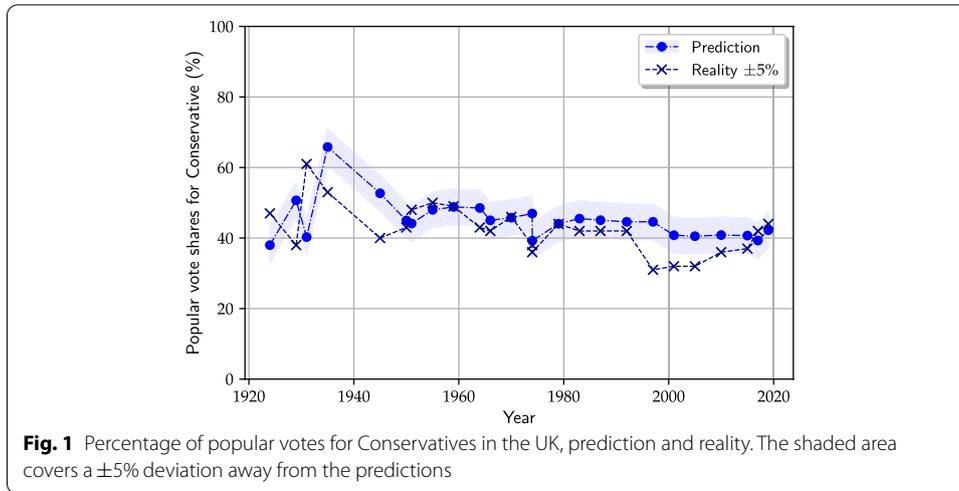
Left: United Kingdom. Right: United States

$n = 100$, the number of possible couples (s_0, s_1) is small enough here that (7) can be solved by directly computing the sum for each of these couples individually. The optimal value s_1^* for s_1 then gives us an estimation of the percentage of votes “locked” by the Conservative party, proportion of the population that will always root for them. The optimal value s_0^* for s_0 is an estimate of the quantity of such votes for all other parties aggregated.

To make a forecast for the i th election, we just have to apply Theorem 1 with $Q = Q^{(s_0^*, s_1^*)}$, $n_1 = x_{i-1}$ and $t = t_i - t_{i-1}$. Equation 4 then gives us the expected percentage \tilde{x}_i of votes gathered by Conservatives on that occasion. This can then be compared to the actual value x_i to assess the efficacy of our approach.

Results for the UK

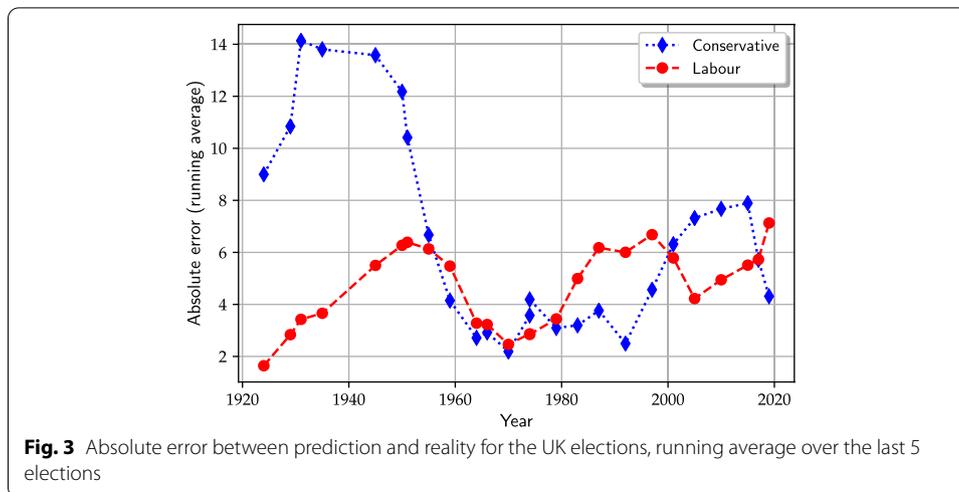
We show in Table 1 (left) the estimated values for (s_0^*, s_1^*) , updated with each new election. They seem to globally stabilise between 15 and 25 for both parties. Look at the last value in the Labour case for example, which is (24, 15). According to our model, this means there is an estimated proportion of 15% of voters that will *always* vote Labour and 24% that will *never* do so. Note that these estimates fluctuate according to the variability of the data. For example in 1922 and 1923 there were twice in a row 38% votes for



Conservative⁴ and as a result it was estimated that 38% of individuals will *always* vote Conservative and the remaining 62% *never* will. This is indeed what maximises the likelihood, with this configuration yielding a probability of 1 for the observed values. On the other hand, with pro-Conservative votes jumping from 38 to 61% in 1935, estimated values of s_0 and s_1 dropped significantly to account for the wide range covered by the data.

In Figs. 1 and 2 we compare our predictions, that is the expectations \tilde{x}_i , with the real outcomes x_i . We plot both values for each election starting with the third one that took place in 1924, because the optimisation problem (7) requires $i \geq 3$. For both parties, most values seem to fluctuate around the 40% mark. The global tendency of the real outcomes looks respected by the predictions, albeit with less variability. Also note that most predictions appear to be within a $\pm 5\%$ vicinity of the real values.

⁴ Remember that those value are rounded to the nearest integer to fit the needs of our model—the actual results were 38.5% and 38%.

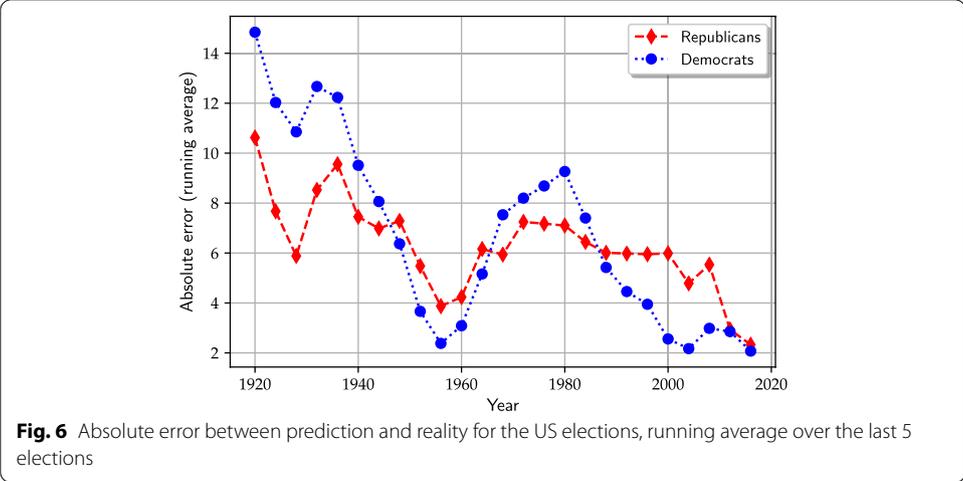
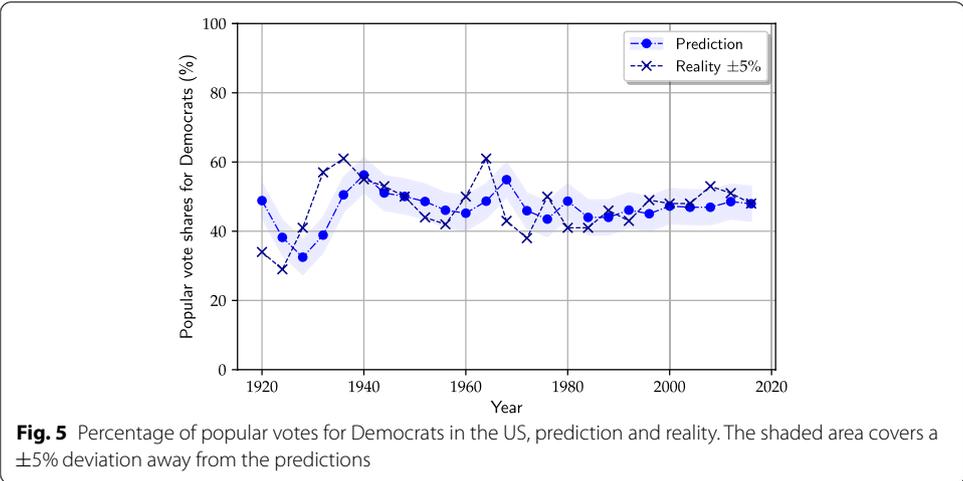
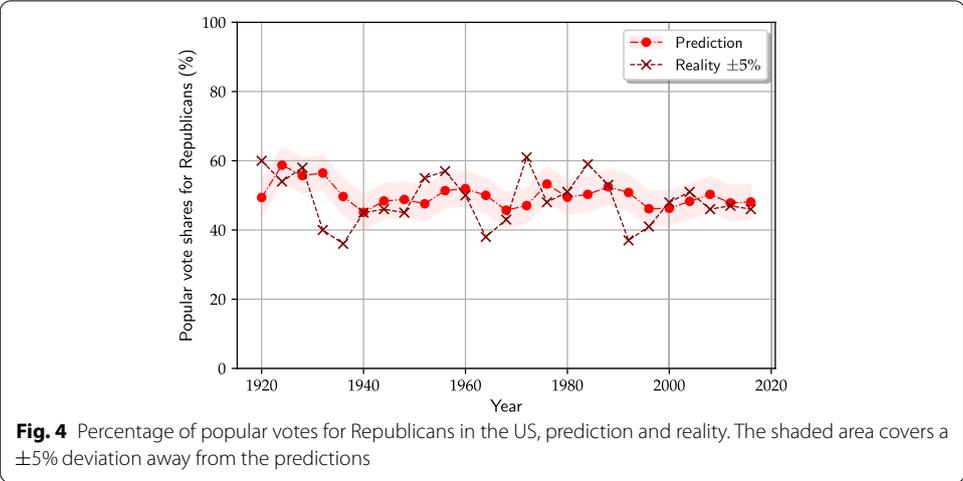


To get a better insight we look at the absolute errors $|\tilde{x}_i - x_i|$ of our predictions. We plot running averages over the last 5 elections in Fig. 3. After a few erratic first years they seem to stabilise between 2 and 8%. More precisely, if we discard the first few years up until 1960 where the model lacks sufficient amount of data to properly calibrate, we get MAEs of respectively 4.63% and 5.23% for Conservative and Labour. Minimal values of 0.06% for Conservatives in 1979 and 0.40% for Labour in 2001 are observed, showing that our method was able to make very accurate predictions in these cases. Surprisingly however, the errors do not seem to monotonically decrease over time, but rather fluctuate. As a matter of facts, peak absolute errors were observed in 1983 (Labour, 13.0%) and 1997 (Conservative, 13.6%).

Results for the US

We apply the exact method described above to the case of presidential elections in the United States. As we did before we independently consider two cases, Republicans versus non-Republicans and Democrats versus non-Democrats. Presidential elections in the US take place every 4 years and we start with the year 1912, then 1916, 1920, and so on. Here again, keep in mind that due to how the American system work, the party with the most popular votes does not necessarily win the elections. The first estimation we are able to make is based on the first two elections and thus our first prediction is for 1920.

We observe similar results as in the UK case. Stubborn values (Table 1, right) estimated (s_0^* , s_1^*) are close, albeit a little bit lower—stabilising at (18, 17) for Republicans and (16, 14) for Democrats. Regarding the predictions (Figs. 4, 5) we again see a majority of them within a 5% margin from the actual outcomes, and a prediction curve that looks more stable than the slightly spiky ones with real values. Note that because of the two-party system in place in the United States, both Republicans and Democrats see their share of popular votes fluctuate around the 50% mark. In the previous case, it was rather around 40% because of the space occupied by smaller parties such as Liberal Democrats or Scottish National Party amongst others. The two-sided aspect of our model—always



one party (0) versus another (1)—may thus be more adapted to the study of the US system.

As for the errors, running averages over the last 5 elections are shown in Fig. 6. Here again after a few erratic first years values appear to be comprised between 2 and 8%. However, where errors in the UK case seemed to increase in the last few years, here they to are dropping down. In fact, our most accurate forecast regarding Democrat votes is for 2016, with only 0.04% error. For Republicans it is in 1940 with 0.10%. Peak errors were again around 13% for both parties, in 1972 (Republicans, 14.0%) and 1964 (Democrats, 12.3%). The MAE over all elections, starting in 1940 when forecasts start to stabilise, is 4.27% for Republicans and 4.83% for Democrats. This is slightly better than in the UK case (4.63% and 5.23%). The MAE error over both cases is then 4.74%.

Conclusion and future work

In this paper we proposed a new method for the forecast of elections results. A lot of published work have used Twitter data for this purpose, usually applying machine learning algorithm to extract sentiment from tweets and estimate a candidate's popularity this way. Despite promising results, such methods have been criticised in the past few years, with problem ranging from bot presence to text mining reliability that cast doubt over their reliability. As such, our model does not rely on Twitter data at all. Instead, we used official results of past elections in the United Kingdom and in the United States to try and predict outcomes of future ones.

Our method is based on findings from a previous work of ours, where we conducted a theoretical analysis of the voter model with stubborn on strongly-connected graphs. Here we applied those in the case to try and predict the percentage of popular votes won by Conservative and Labour parties in the United Kingdom, and the percentage of popular votes collected by the Republican and Democratic parties in the United States. To do so, we considered official results of past elections as observations of independent realisations of the voter model. From there we were able to perform time-evolving estimates of the model parameters and use them to forecast an outcome.

Our model yielded an MAE of 4.74%, reaching absolute errors as low as 0.04% and as high as 14%. In their review, (Gayo-Avello 2013) suggest that any model used to predict the elections outcome should not have an MAE higher than 1 or 2%. This is because the result of an election is more often than not the matter of just a few percents. According to this standard, our MAE is not low enough to reliably predict the outcome of an election. Some previous works reached error averages as low as 0.63 (Saleiro et al. 2016) and 1.65% (Tumasjan et al. 2011). Additionally, we tested our method against the baseline of systematically predict the exact result of the previous election. This simple method returned an overall 5.03% average error, which is not much worse than the 4.74% obtained via our method. Moreover, the first few elections results were discarded in both cases, as it was deemed that the model did not have enough data at this point to make predictions with a high enough confidence. The choice of a limit though is made on the basis of our observation of the model's behaviour and is purely subjective. Changing the limit would in turn make for different results that might be better or worse.

Although our method did not yield significant enough results here, we believe it is an interesting step in a novel direction. The use of results from previous elections provides

a new take on the matter, which only relies on official data. Also our model does not only forecast the elections results, it also gives us estimates of the proportion stubborn voters, that is the proportion of individuals who will *always*—or *never*—vote for the considered parties. This provides meaningful insight on the political landscape of the considered areas.

Several extensions of the model could be considered to improve its accuracy. First of all, adding in-between election polls to the data would go a long way in improving the estimates. With a few years gap from one election to another, it is too wide a range of possibilities for the model to account for. Second, one could take a deeper look into the past of a country's results and try to detect tendencies about landslide victories, incumbency reelection and so forth. We believe that having a deeper understanding of the specific country one is working with could substantially improve the model calibration process. Finally, combining our method with Twitter data-based estimations may lead to higher accuracy.

Abbreviations

MAE: Mean Absolute Error; UK: United Kingdom; US: United States.

Authors' information

Antoine Vendeville is a PhD student at University College London (United Kingdom) in the Computer Science department. Benjamin Guedj is a Principal Research Fellow in machine learning at University College London and a tenured research scientist at Inria, France. Shi Zhou is an Associate Professor at the Department of Computer Science, University College London. All three authors are affiliated with UCL's Centre for Doctoral Training in Cybersecurity. All three authors are affiliated with UCL's Centre for Artificial Intelligence.

Acknowledgements

Not applicable.

Author's contributions

This is a joint work by the three authors, with A.V. being the leader of the project. All authors read and approved the final manuscript.

Funding

This project was funded by the UK EPSRC grant EP/S022503/1 that supports the Centre for Doctoral Training in Cybersecurity delivered by UCL's Departments of Computer Science, Security and Crime Science, and Science, Technology, Engineering and Public Policy.

Availability of data and materials

All code used is available online at <https://github.com/AntoineVendeville/HowOpinionsCrystallise>. Data for the United Kingdoms elections is available online (Audickas et al. 2020). Data for the United States elections has been crawled from Wikipedia (https://en.wikipedia.org/wiki/United_States_presidential_election#Popular_vote_results).

Competing interests

The authors declare that they have no competing interests.

Author details

¹ Department of Computer Science, University College London, London, UK. ² Centre for Doctoral Training in Cybersecurity, University College London, London, UK. ³ Centre for Artificial Intelligence, University College London, London, UK. ⁴ Inria Lille - Nord Europe Research Centre, Lille, France.

Received: 22 September 2020 Accepted: 1 December 2020

Published online: 07 January 2021

References

- Audickas L, Cracknell R, Loft P (2020) UK election statistics: 1918-2019—a century of elections. <https://commonslibrary.parliament.uk/research-briefings/cbp-7529/>
- Bovet A, Morone F, Makse HA (2018) Validation of Twitter opinion trends with national polling aggregates: Hillary Clinton vs Donald Trump. *Sci Rep* 8:8673. <https://doi.org/10.1038/s41598-018-26951-y>
- Caldarelli G, Chessa A, Pammolli F, Pompa G, Puliga M, Riccaboni M, Riotta G (2014) A multi-level geographical study of Italian political elections from twitter data. *PLoS ONE* 9(5):1–11. <https://doi.org/10.1371/journal.pone.0095809>

- Chen P, Redner S (2005) Majority rule dynamics in finite dimensions. *Phys Rev E*. <https://doi.org/10.1103/PhysRevE.71.036101>
- Clifford P, Sudbury A (1973) A model for spatial conflict. *Biometrika* 60(3):581–588. <https://doi.org/10.1093/biomet/60.3.581>
- Dokoohaki N, Zikou F, Gillblad D, Matskin M (2015) Predicting Swedish elections with twitter: a case for stochastic link structure analysis. In: IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM), pp 1269–1276. <https://doi.org/10.1145/2808797.2808915>
- Fernley J, Ortgieste M (2019) Voter models on subcritical inhomogeneous random graphs. [arXiv:1911.13187](https://arxiv.org/abs/1911.13187)
- Fink C, Bos N, Perrone A, Liu E, Kopecky J (2013) Twitter, public opinion, and the 2011 Nigerian presidential election. In: International conference on social computing, pp 311–320. <https://doi.org/10.1109/SocialCom.2013.50>
- Garcia ACB, Silva W, Correia L (2018) The PredNews forecasting model. In: Proceedings of the 19th annual international conference on digital government research: governance in the data age. Association for Computing Machinery, New York. <https://doi.org/10.1145/3209281.3209295>
- Gayo-Avello D (2012) No, you cannot predict elections with twitter. *IEEE Internet Comput* 16(6):91–94. <https://doi.org/10.1109/MIC.2012.137>
- Gayo-Avello D (2013) A meta-analysis of state-of-the-art electoral prediction from twitter data. *Soc Sci Comput Rev* 31(6):649–679. <https://doi.org/10.1177/0894439313493979>
- Grimaldi D, Cely JD, Arboleda H (2020) Inferring the votes in a new political landscape: the case of the 2019 Spanish presidential elections. *J Big Data*. <https://doi.org/10.1186/s40537-020-00334-5>
- Hassin Y, Peleg D (2002) Distributed probabilistic polling and applications to proportionate agreement. *Inf Comput* 171(2):248–268. <https://doi.org/10.1006/inco.2001.3088>
- Holley RA, Liggett TM (1975) Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann Probab* 3(4):643–663. <https://doi.org/10.1214/aop/1176996306>
- Huberty ME (2013) Multi-cycle forecasting of congressional elections with social media. In: Proceedings of the 2nd workshop on politics, elections and data. PLEAD '13. Association for Computing Machinery, New York, pp 23–30. <https://doi.org/10.1145/2508436.2508439>
- Jungherr A, Jürgens P, Schoen H (2012) Why the pirate party won the German election of 2009 or the trouble with predictions: a response to Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. "Predicting elections with twitter: What 140 characters reveal about political sentiment". *Soc Sci Comput Rev* 30(2):229–234. <https://doi.org/10.1177/0894439311404119>
- Jungherr A, Schoen H, Posegga O, Jürgens P (2017) Digital trace data in the study of public opinion: an indicator of attention toward politics rather than political support. *Soc Sci Comput Rev* 35(3):336–356. <https://doi.org/10.1177/0894439316631043>
- Klamser PP, Wiedermann M, Donges JF, Donner RV (2017) Zealotry effects on opinion dynamics in the adaptive voter model. *Phys Rev E*. <https://doi.org/10.1103/PhysRevE.96.052315>
- Mobilia M (2003) Does a single zealot affect an infinite group of voters? *Phys Rev Lett* 91:028701. <https://doi.org/10.1103/PhysRevLett.91.028701>
- Mobilia M, Petersen A, Redner S (2007) On the role of zealotry in the voter model. *J Stat Mech Theory Exp* 2007(08):08029–08029. <https://doi.org/10.1088/1742-5468/2007/08/p08029>
- Mossel E, Neeman J, Tamuz O (2014) Majority dynamics and aggregation of information in social networks. *Auton Agent Multi Agent Syst* 28(3):408–429. <https://doi.org/10.1007/s10458-013-9230-4>
- Mukhopadhyay A, Mazumdar RR, Roy R (2020) Voter and majority dynamics with biased and stubborn agents. *J Stat Phys*. <https://doi.org/10.1007/s10955-020-02625-w>
- O'Connor BT, Balasubramanyan R, Routledge BR, Smith NA (2010) From tweets to polls: Linking text sentiment to public opinion time series. In: ICWSM
- Perron E, Vasudevan D, Vojnovic M (2009) Using three states for binary consensus on complete graphs. *IEEE INFOCOM* 2009:2527–2535. <https://doi.org/10.1109/INFCOM.2009.5062181>
- Phillips L, Dowling C, Shaffer K, Hodas NO, Volkova S (2017) Using social media to predict the future: a systematic literature review. *CoRR* [abs/1706.06134](https://arxiv.org/abs/1706.06134)
- Saleiro P, Gomes L, Soares C (2016) Sentiment aggregate functions for political opinion polling using microblog streams. In: Proceedings of the ninth international conference on computer science and software engineering. C3S2E '16. Association for Computing Machinery, New York, pp 44–50. <https://doi.org/10.1145/2948992.2949022>
- Sood V, Tibor A, Redner S (2008) Voter models on heterogeneous networks. *Phys Rev E* 77:041121. <https://doi.org/10.1103/PhysRevE.77.041121>
- Thapen NA, Ghanem MM (2013) Towards passive political opinion polling using twitter. *CEUR Workshop Proc* 1110:19–34
- Tumasjan A, Sprenger TO, Sandner PG, Welpe IM (2011) Election forecasts with twitter: how 140 characters reflect the political landscape. *Soc Sci Comput Rev* 29(4):402–418. <https://doi.org/10.1177/0894439310386557>
- Usher J, Dondio P (2020) Brexit election: forecasting a conservative party victory through the pound using arima and facebook's prophet. In: Proceedings of the 10th international conference on web intelligence, mining and semantics. WIMS 2020. Association for Computing Machinery, New York, pp 123–128. <https://doi.org/10.1145/3405962.3405967>
- Vendeville A, Guedj B, Zhou S (2020) Voter model with stubborn agents on strongly connected social networks. [arXiv:2006.07265](https://arxiv.org/abs/2006.07265)
- Yildiz ME, Pagliari R, Ozdaglar A, Scaglione A (2010) Voting models in random networks. In: Information theory and applications workshop (ITA), pp 1–7. <https://doi.org/10.1109/ITA.2010.5454090>
- Yildiz ME, Ozdaglar A, Acemoglu D, Saberi A, Scaglione A (2013) Binary opinion dynamics with stubborn agents. *ACM Trans Econ Comput*. <https://doi.org/10.1145/2538508>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.