Considering Opinion Dynamics and Community Structure in Complex Networks:
A view towards modelling elections and gerrymandering

James F. Wall
University of Oxford
September 2008

Electoral districting and gerrymandering are problems as old at the United States itself. In this project, by employing opinion dynamics on some social networks and using their community structure, we seek to model elections in a two party political system like the USA and simulate the phenomenon that gerrymanders attempt to use to their advantage - electorates within different districts or spheres of influence tend to vote for different parties. We show that this effect can be modelled successfully on a network by the voter model if community structure is strong, or by slightly altering the dynamics of the voter model to include a feedback mechanism.

INTRODUCTION

The *network* is a concept that is ubiquitous in physical and human sciences; a network is a configuration of agents (*vertices* or *nodes*), together with connections between them (*edges*) which represent some manner of interaction. Examples include biological networks such as food webs, where the vertices and edges denote organisms and consumption respectively; information networks such as the World Wide Web, in which the vertices are web pages and the edges hyperlinks; and social networks where the nodes may correspond to people and the edges friendship [1].

It is partly due to this universality that network theory has become a burgeoning field of study in recent years; more influential has been the ability, granted by the availability of greater computer power, to collate the data that forms such large networks, sometimes of millions even billions of nodes, and to conduct statistical analysis of their structure [1, 2].

Part of this research has involved modelling on social networks collective patterns of behaviour that emerge from individual interactions; these social phenomena include opinion, cultural and language dynamics [3]. For instance, one may simulate the spreading of support for a political party within an electorate [4]: the goal is to find some simple rules of exchange between the nodes, each of which is endowed with its own opinion, that propagates said opinion and reproduces patterns observed in real democratic systems.

One of democracy's longest running points of contention and forms of electoral skulduggery is the *gerrymander*. Gerrymandering refers to the practice of drawing electoral district lines in order to advantage a particular political party or parties [5]. This can and has taken several forms, for instance there is partisan gerrymandering where districts are redefined in order to favour one party over another; bipartisan, or 'sweetheart', gerrymandering is where parties collude in redistricting in order to secure the re-election of incumbents; and there is gerrymandering along ethnic-minority lines so that, depending on the motivation, the vote of minorities is either diluted or strengthened [6].

The term gerrymander itself originates from 1812 when Massachusetts governor, Elbridge Gerry, presided over his party's strategic redrafting of a long, sinuous electoral district with the aim of winning the consequent senatorial election by reducing competition from their opponents. This district was likened in print to a salamander, and hence the contemporary neologism *gerrymander* was coined from the concatenation of the governor's surname and -*mander*.

Electoral districting remains an active debate and area of study today, its problems and possible solutions have been discussed not only in political science [6, 7, 8] but also in mathematics [9, 10], computer science [11, 12], physics [5], and law. Yet, the focus of the majority of literature remains concerned with the nation whose legislature gave birth to the gerrymander - the United States. This is perhaps the case because the United States' status as the world's superpower puts it in a unique position of scrutiny a propos its constitution and the opportunity of gerrymandering it affords. After every decennial census the House of Representatives is reapportioned in accordance with the population changes observed, with each state of the Union granted a number of seats proportional to its population. Say a state is apportioned n seats in the House; the state legislature is then required to divide its territory into n electoral districts of equal population which will each elect one representative to take a seat in Congress. It is the power granted to the incumbent state authorities to redraw these boundaries, rather than to an independent body such as the courts or an electoral commission, that exposes American voters to this electoral mischief.

OVERVIEW

In this paper, we consider the results obtained by simulating elections on networks acquired from the social networking site *facebook.com*. Members of the site create a personal profile and can add other members as 'friends', creating links between their own and their friends' profiles, hence members and these 'friendships' form the vertices and edges of a readily accessible social network. At the time the data was collected users of facebook required a valid university email address and so the networks we consider consist of exclusively students at the same institution [13]. The individual college networks we use in this paper are those of the California Institute of Technology (CA), commonly abbreviated to Caltech, and two liberal arts colleges Reed (OR) and Haverford (PA); these networks were chosen primarily for their small size for ease and speed of computation.

On these networks we ran variants upon a particular opinion dynamic in order to model an 'election'. In reality of course, most elections conducted in public life - at least the ones in which the effects of gerrymandering has been considered - involve far larger populations than those in our facebook networks; instead our subject networks aim to act effectively as a sampling of a population, or as a surrogate for an electorate.

It is known from our everyday experience and from studies conducted that individuals tend to associate with others who are like themselves, for example in terms of race, age or income [2]. This is no different in political opinion or political party affiliation: consider the population of a region or nation at large, we expect it be more likely that rural communities will vote Republican and urban communities Democrat (or Conservative-Labour if we cross the Atlantic). On a smaller scale, we expect within a friendship group similar political leanings just as we expect similarity in other cultural, racial or economic traits (all four of course being complicatedly interconnected). This clustering of votes for particular political parties on the grounds of neighbourhood or friendship group is what we seek to model in this paper; furthermore, we investigate if these clusters can be harnessed so that these social networks can be segregated in a natural way as to influence the result of a possible election in the same manner a gerrymander aims to.

MODELLING ELECTIONS

We model opinion as a binary variable: several such models have been studied on networks for example the voter, the Sznadj and majority-rule models [3], but it is the former [14] and variations upon it that we shall consider on our college networks. In a binary variable dynamic every node is in one of two states, often termed the *spin* of a node thanks to the study of thid field in statistical physics with regard to ferromagnets [3]. In our case we form a two party electoral system: one state may be 'vote Republican' the other 'vote Democrat' (assuming everybody votes). In all our discussions we take the initial configuration of votes on the network to be disordered: every node has an equal probability of voting Republican or Democrat. It is generally assumed that initially heterogeneity dominates [3], without interaction with others every individual expresses a purely personal response on the question of which party they support (of course one could argue that the formation of friendships, which are the edges of our network, involve some prior interaction between the nodes anyway).

In real world networks typically the distribution of edges is both globally and locally inhomogeneous [15]. One observes a high concentration of edges within certain groups of nodes and a low number of edges between these groups, these mesoscopic features of a network are called *communities* [2]. In social networks they may correspond to friendship groups, or in the World Wide Web to groups of web pages on related topics. Recently a variety of algorithms have been designed to carry out community detection; in this paper we employ one of these computational techniques - the Newman's leading eigenvector method [16].

We shall view the communities detected by the eigenvector method on our networks to be analogous to electoral districts, and consider the votes for the two respective parties within these communities over time and as we vary the voter model. Clearly we cannot equate community structure with real electoral districts - districts are not drawn along lines of acquaintance but by the geography of the region, our networks do not consider the physical proximity of those who constitute the nodes only the notion of 'friendship' which, now in times of globalization and greater mobility, no longer necessarily implies close proximity of habitation. However, considering the votes within each community enables us to investigate the behaviour of our opinion dynamics with regard to the mesoscopic structure of our networks, and in particular the aforementioned social phenomenon that people within the same community share similar voting habits.

THE VOTER MODEL

The voter model is possibly the simplest opinion dynamic and hence has been studied in great detail both on regular lattices and real world networks; it is one of the few non-equilibrium stochastic processes that can be solved in any dimension [3]. All nodes are endowed with some spin $\sigma = \pm 1$(upspin or downspin); a node i is chosen at random, then one of its neighbours, j say, is randomly chosen, the former then assumes the opinion of the latter, i.e. $\sigma_i = \sigma_j$. In this paper, we term a single update to be when the spin of single node is changed from one

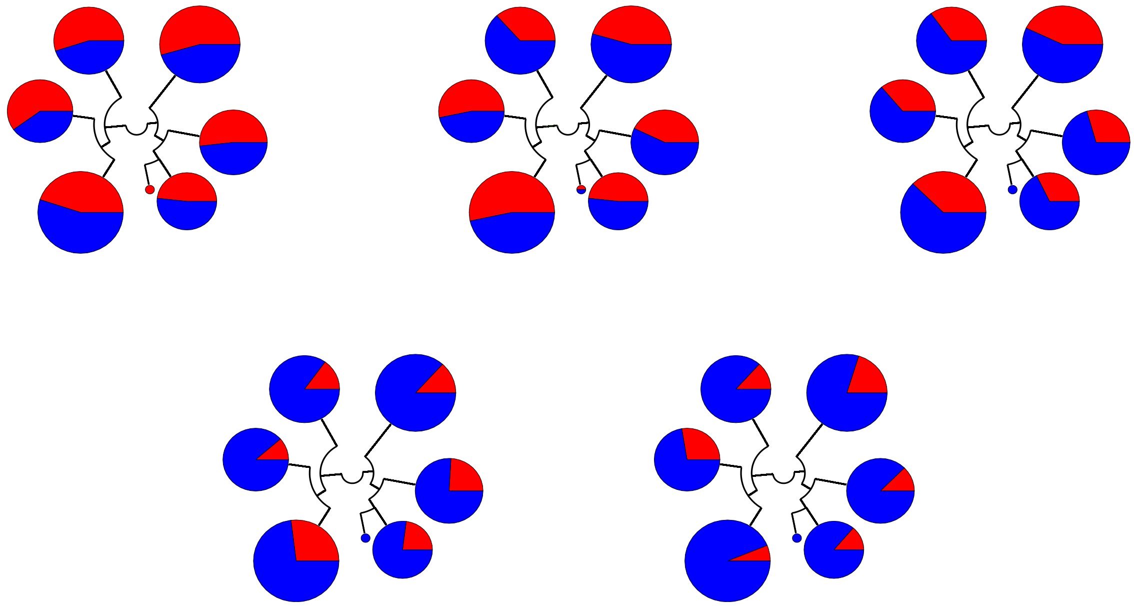

Figure 1: Pie chart dendrograms for a run of the voter model on the Caltech network (a) initially, and after (b) 5,000, (c) 10,000, (d) 15,000, (e) 20,000 updates. The 7 pies correspond to the 7 communities found when the leading eigenvector method is employed on the Caltech network; the radius of each pie is determined by the number of nodes in the corresponding community. The pies themselves are colour coded according to vote or spin: blue signifying upspin (or 'vote Democrat') and red downspin (or 'vote Republican'). For more details see [13]

state to another, if two nodes of the same opinion are chosen we re-choose until two nodes of different spin are picked (some other papers choose to define their 'updates' in a different manner).

The voter model is motivated by the idea of 'social influence' - that individuals imitate their neighbours and hence people tend to become more alike, whilst the state of the population as a whole does not play a direct role [3]. Indeed in any network with a finite number of nodes, such as ours, the dynamics always eventually reach consensus - either all nodes with up spin or all with downspin [14].

The Caltech network consists of 762 nodes, and employing the leading eigenvector method yields 7 communities; one community (District 4) contained only two nodes and thus was neglected from our analysis. We ran the voter model from a disordered initial configuration of party votes several times, each time up to consensus or to 20,000 updates, whichever occurred first. We observed that the proportion of votes for either party locally within the districts largely reflected the party's global vote. Figure 2 is a line plot of the proportion of the vote in the different communities over one run of the voter model: it illustrates that the division of the vote between the two parties does not vary greatly across the 6 main communities or 'districts' as the number of updates of the model increases. Figure 1 provides a second visualization of the same run of the voter model - it shows 5 pie chart dendrograms which display the proportion of the vote in the 7 communities of the network in its initial disordered state and after 5,000, 10,000, 15,000 and 20,000 updates. The dendrograms' pies correspond to the communities yielded by the leading eigenvector method, and the radius of each pie is determined by the number of nodes in the corresponding community. The pies themselves are colour coded according to vote or spin: blue signifying upspin (or 'vote Democrat') and red downspin (or 'vote Republican'). It is observed that all the pies in a single dendrogram are relatively similar, in accordance with what I have explained.

Runs of the voter model were also conducted on the 962 and 1445 node Reed and Haverford networks, for which the leading eigenvector method returns 4 and 8 communities respectively (again small communities, namely Districts 7 and 8 of the Haverford network are omitted from analysis). Similar behaviour to that described for the Caltech network was observed, yet notable anomalies did occur. Again, runs of the voter model suggested that over time the proportions of the vote in most districts of the same network were similar; however, it was observed that typically the proportions of the vote in District 2 of the Reed network and District 5 of the Haverford network often differed noticably from that in other districts of the same network. Line plots Figures 3 and 4 provide a visualization of this - Figure 3 represents a run of the voter model on the Reed

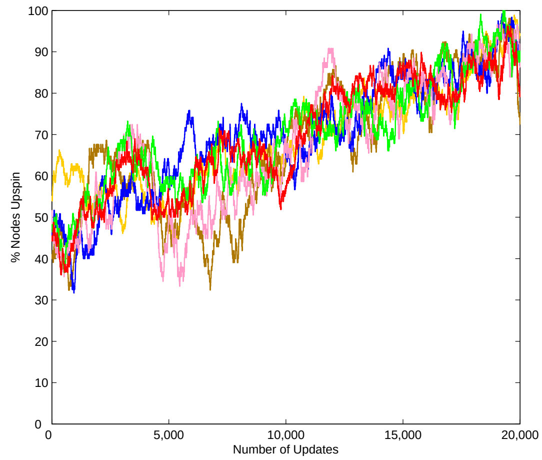

Figure 2: [Colour] A run of the voter model on the Caltech network. Each line expresses the proportion of upspin nodes (or votes for the Democrats) in one of the communities. A line showing the proportions of the vote in the two node District 4 is omitted for clarity.

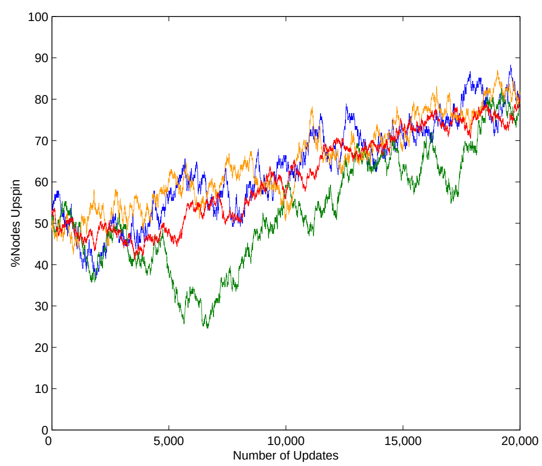

Figure 3: [Colour] A run of the voter model on the Reed network. Each line expresses the percentage of upspin nodes in a particular community. The green line corresponds to District 2.

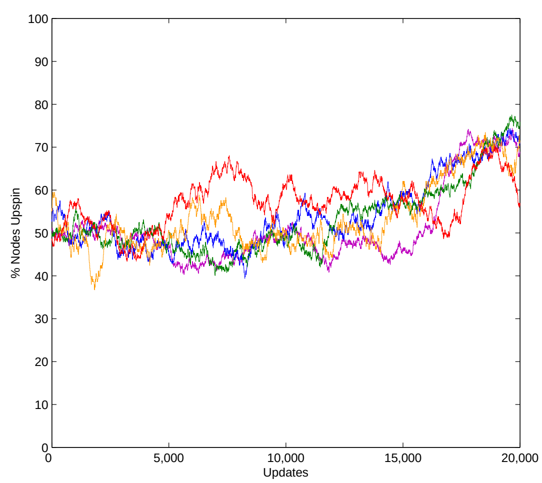

Figure 4: [Colour] A run of the voter model on the Haverford network. Each line expresses the percentage of upspin nodes in a particular community. The red line corresponds to District 5.

network (the green line representing the vote in District 2), and Figure 4 a run on the Haverford network (the red line representing the vote in District 5).

We undertook further analysis of the communities of our 3 networks to account for the differences observed across the districts when the voter model was run. Community detection algorithms like the eigenvector method partition networks so that every node is binned into a community; intuitively we expect that some of these communities to be 'stronger' or 'better defined' than others. How can we pin down and quantify this vague notion of how well defined or isolated individual communities are?

The leading eigenvector method depends on maximizing a quality function called modularity - a global measure of the quality of a partition [16]. We utilized instead a measure of local modularity - a quantity associated with a single community or module rather than a partition of communities [17]. Suppose **C** is a community, consider **B** - the set of vertices that are on the boundary of **C** (i.e. nodes that have at least one neighbour which is not in **C**). Intuitively, we think of **C** having a sharp boundary if the proportion of edges which have a node in **B** at one end and a node in **C** at the other is much greater than proportion of edges which have a node in **B** at one end and have a node that is not in **C** at the other [17]. It is this idea of a sharp boundary that the measure of local modularity takes to be the defining factor of a well defined or isolated community.

The boundary adjacency matrix B_{ij} of community **C** is defined to be:

$$B_{ij} = \begin{cases} 1 & 1 \text{ if nodes } i \text{ and } j \text{ are connected and} \\ & \text{at least one of them is in } \mathbf{B} \\ 0 & \text{otherwise} \end{cases}$$

The local modularity, R, of community **C** is defined to be:

$$R = \frac{\sum B_{ij}\,\delta(i,j)}{\sum B_{ij}} = \frac{I}{T}$$

where $\delta(i,j)$ equals 1 if either node i is in **C** and node j is in **B** (or vice versa) and 0 otherwise, and where T is the number of edges with one or more endpoints in **B**, and I is the number of edges which have both endpoints in **C** and at least one of them is in **B** [17]. The local modularity R takes a value between 0 and 1 which is proportional to the sharpness of the boundary **B** of community **C**. We calculated local modularity for each of the respective communities in the Caltech, Reed and Haverford networks, see Table 1. All the Caltech communities have local modularities below 0.5, as do the majority of districts of the Haverford and Reed networks. However, District 2 of the Reed and 5 of the Haverford networks have a relatively high local modularity in comparison; so we may speculate that how isolated or distinct these

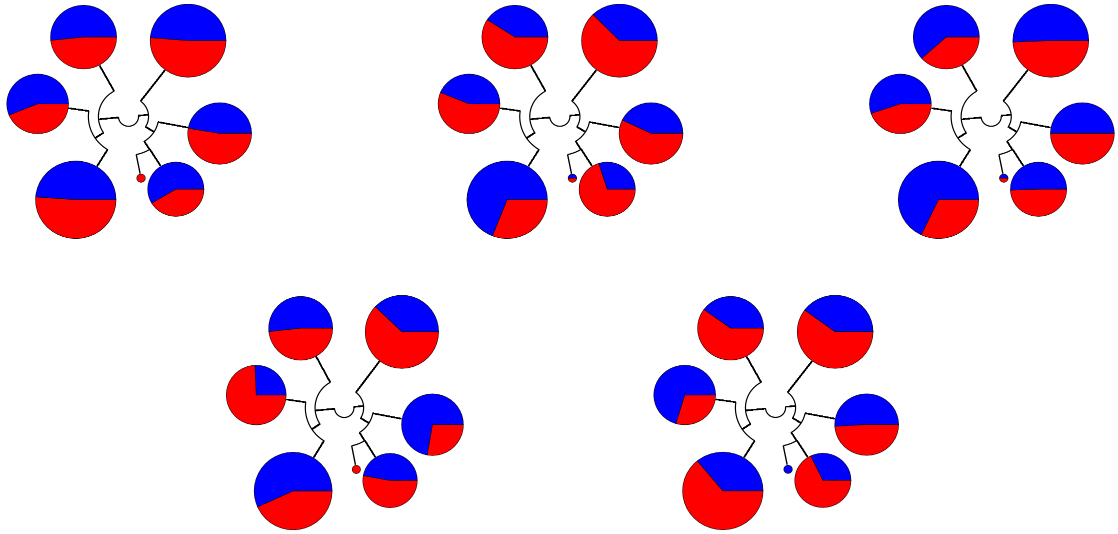

Figure 5: Pie chart dendrograms for a run of the voter model on the modified Caltech network where the communities are complete subgraphs (a) initially, and after (b) 5,000, (c) 10,000, (d) 15,000, (e) 20,000 updates.

Network & District	No. of Nodes	R
Caltech 1	158	0.4250
Caltech 2	112	0.3436
Caltech 3	87	0.3722
Caltech 5	178	0.3678
Caltech 6	105	0.3436
Caltech 7	120	0.3024
Reed 1	128	0.2683
Reed 2	179	0.5301
Reed 3	184	0.3345
Reed 4	471	0.5957
Haverford 1	434	0.3763
Haverford 2	226	0.2588
Haverford 3	307	0.3273
Haverfoed 4	158	0.1831
Haverford 5	288	0.6353
Haverford 6	25	0.1507
Caltech* 1	158	0.8094
Caltech* 2	112	0.7409
Caltech* 3	87	0.6227
Caltech* 5	178	0.8369
Caltech* 6	105	0.7076
Caltech* 7	120	0.7094

Table 1: Local Modularity of Communities, Caltech* denotes the modified version of the Caltech network where each community is a complete subgraph.

communities are has some bearing on the voting behaviour observed.

To try and establish this connection more fully we asked how the voter model would behave if we artificially strengthened the community structure of the Caltech network by making each of the 7 communities a complete subgraph (i.e. adding edges to the network so that every node is connected to every other node in its community). Table 1 displays the local modularities of our new communities, the values of which are considerably higher than the corresponding communities in the original Caltech network and serve as testament to the fact that these are much more clearly distinct districts. A typical run of the voter model on this new network is illustrated in Figures 5 and 6. Figure 5 shows on this new artificial network there are enormous variations in the voting proportions across communities, in comparison to the original Caltech network of Figure 2. As a result of these greatly fluctuating voting proportions over time, if we view our communities as electoral districts we have different 'winners' in different districts over time. In the former run of the voter model since the proportion of votes in each district was similar over time each returned the same winner; however as illustrated in Figure 6 the current network returns a plurality of winners over time: in some districts the Republicans win and in others the Democrats win.

Communities with high local modularity have a lower proportion of edges with end points in other communities, hence if the initial randomly selected node the voter model chooses is in such a community the randomly chosen neighbour involved in the subsequent dynamic is more likely to belong to it also. Thus the stronger the community structure of a network the greater degree of independence in the voting dynamics of each community, resulting in visibly differing voting proportions over time. We should note however that the communities of the original Caltech network reflect the House structure of the institution: the fact that the original communities are in some sense 'correct' has led to the network being suggested as a benchmark network to test future methods [13]. From our perspective it also suggests that even though communities correspond roughly to what one would expect we in fact require communities to be very pronounced to display appreciably different voter dynamics.

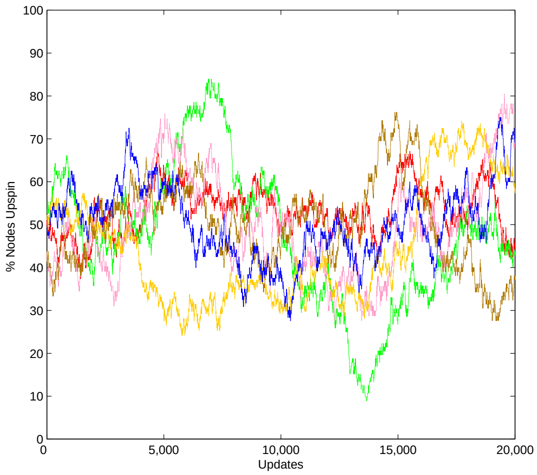

Figure 6: [Colour] A run of the voter model on the modified Caltech network where each community is a complete subgraph. Each line expresses the percentage of upspin nodes in a particular community.

THE VOTER MODEL WITH FEEDBACK MECHANISMS

As mentioned, the classic voter model shall always run to consensus on a network of finite nodes. In the situation of fair multiparty elections however, complete consensus of an electorate is not a realistic outcome - support for more than one party persists. Furthermore, if we take the United States as the template of our two party system the increased polarization of opinion in that nation in recent years leads us to the expectation that the division of the popular vote between our parties should be kept roughly around the 50:50 mark - for instance, examine the popular vote for the two main political parties in the post-war presidential elections (discounting votes for minor third party candidates) displayed in Table 2.

Year	Republican Vote	Democrat Vote
2004	51.49%	48.51%
2000	49.73%	50.27&
1996	45.26&	54.74%
1992	46.54%	53.46&
1988	53.90%	46.10%
1984	59.17%	40.83%
1980	55.30%	44.70%
1976	48.95%	51.05%
1972	61.79%	38.21%
1968	50.40%	49.60%
1964	38.66%	61.34%
1960	49.91%	50.09%
1956	57.76%	42.24%
1952	55.29%	44.71%
1948	47.68%	52.32%

Table 2: Popular vote in the last 15 American presidential elections: we consider only the popular vote of the population who voted either Republican or Democrat, votes for other candidates are omitted.

With this in mind, we altered the dynamic of the voter model to include various feedback mechanisms. Modifications of the voter model which have been the subject of study in the past have included the addition of noise [18]; the so called 'noisy voter model' includes the possibility that a node can flip its spin (or 'change its vote') spontaneously when equal to all its neighbours. In a similar manner we sought to introduce flips of random nodes, the probability of a flip being dependent on the global proportion of up and down-spin nodes, with the aim of swinging the proportion of the vote for our pseudo- Republicans and Democrats about the 50:50 mark over time.

Suppose the voter model is run on a network of N nodes; after every update of the voter model the following dynamic is enacted:

Feedback mechanism FM(i) - Suppose there are p upspin ('vote Democrat') nodes and q downspin ('vote Republican') nodes. If $p \leq q$ then a random downspin node is chosen, and flipped with probability $\frac{q}{N}$. If $p \geq q$ then a random upspin node is chosen, and flipped with probability $\frac{p}{N}$.

For FM(i) consensus is an unstable state, rather than an absorbing state as with the traditional voting model. Note however, that the probability of a flip occurring after an update is always at least 0.5, thus resulting in a high frequency of random flips being imposed on top of the ordinary voter model. In our second feedback mechanism we reduced this number, whilst maintaining consensus to be an unstable state, by making the probability that a flip occurs dependent on the difference between the two votes instead:

Feedback mechanism FM(ii) - Suppose there are p upspin and q downspin nodes. The probability that some node is flipped is $\frac{|p-q|}{N}$. If $p < q$ then this flip will be of some random downspin node to upspin, and if $q < p$ this flip will be of some random upspin node to downspin.

We first investigated runs of FM(i) and FM(ii) on a one dimensional lattice of 100 nodes from an initial disordered state. Figure 7 provides a visualization of a run of the original voter model, FM(i) and FM(ii) on the lattice. We represent the nodes of the lattice at any one time by a strand of cells: each cell corresponds to a node of the lattice and the two neighbouring cells to the adjacent nodes. The colour of a cell expresses the spin of the corresponding node - white for up, black for down. The plots in Figure 7 are created by the strands at each step of the respective models placed one above the other as time goes on. Figure 8 (Left) plots the number of opinion clusters, n_{cl}, of one run of each of the three models

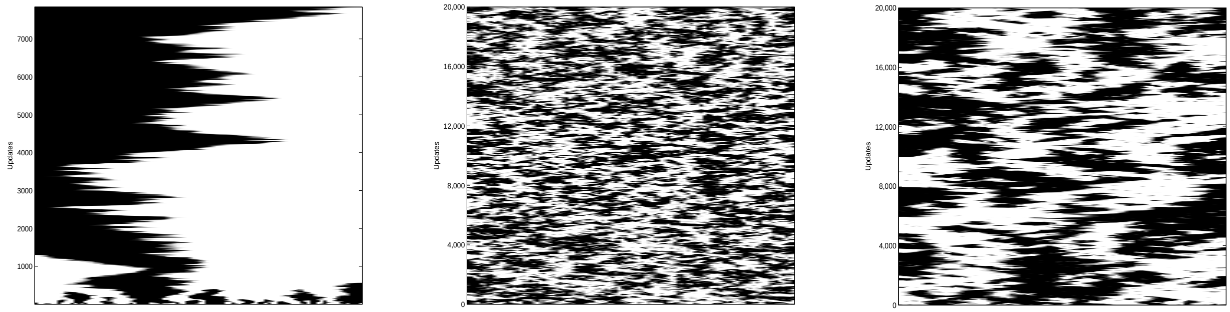

Figure 7: Runs on a 100-node one dimensional lattices of, from left to right, the voter model, FM(i) and FM(ii).

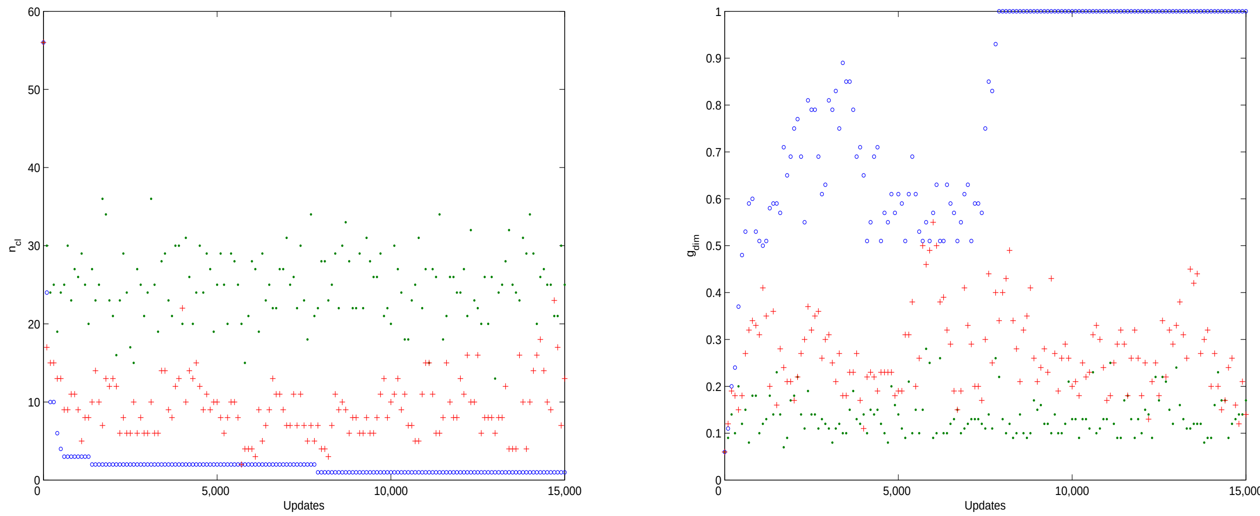

Figure 8: [LEFT] The number of opinion clusters, n_{cl}, and [RIGHT] the relative giant cluster size, g_{dim}, of a run the voter model (blue circles), FM(i) (green dots) and FM(ii) (red crosses) on a 100-node one-dimensional lattice.

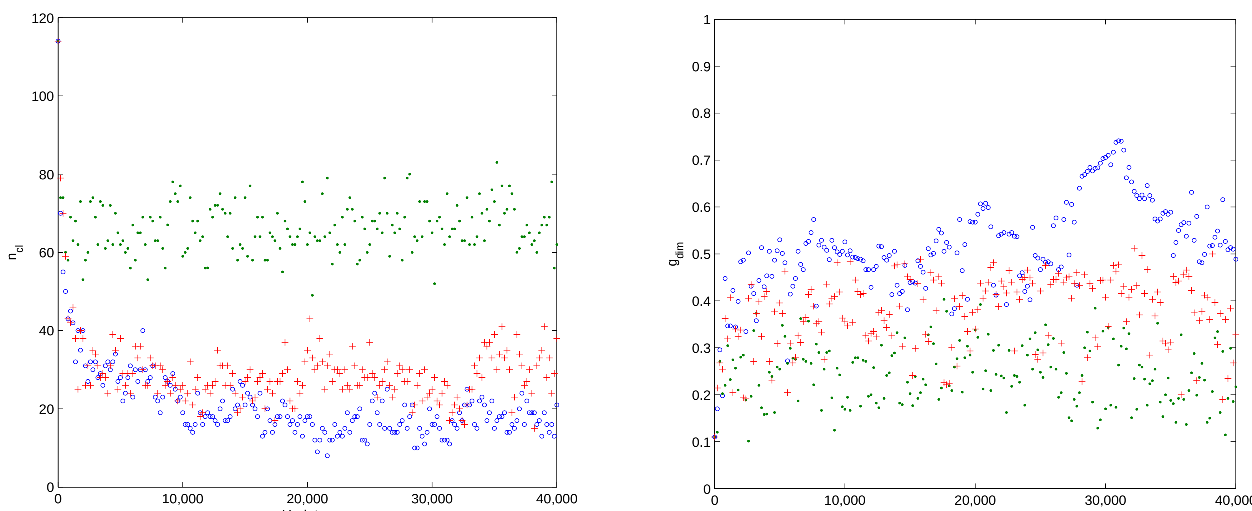

Figure 9: [LEFT] The number of opinion clusters, n_{cl}, and [RIGHT] the relative giant cluster size, g_{dim}, of a run the voter model (blue circles), FM(i) (green dots) and FM(ii) (red crosses) on a 30x30-node two-dimensional lattice.

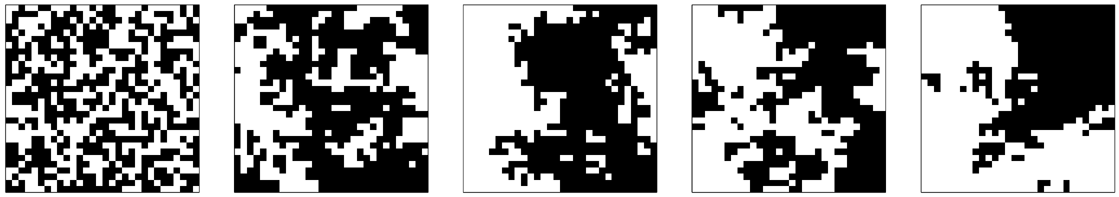

Figure 10: A run of the voter model on a 30x30-node two-dimensional lattice (from left to right): initially, after 5,000, 10,000, 15,000 and 20,000 updates.

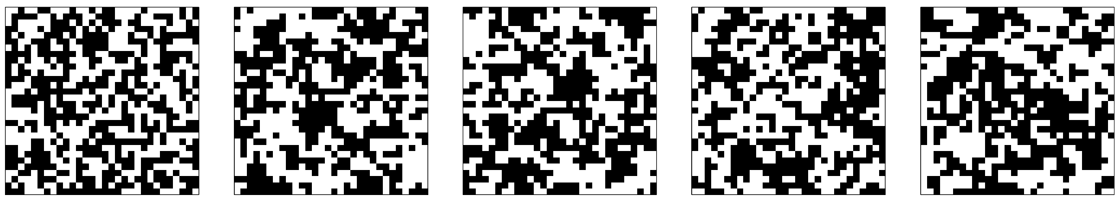

Figure 11: A run of FM(i) on a 30x30-node two-dimensional lattice (from left to right): initially, after 5,000, 10,000, 15,000 and 20,000 updates.

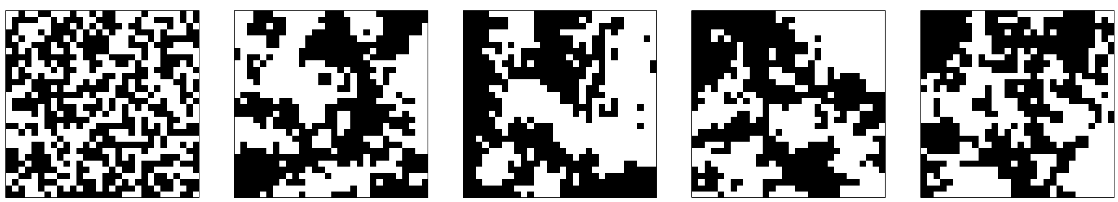

Figure 12: A run of FM(ii) on a 30x30-node two-dimensional lattice (from left to right): initially, after 5,000, 10,000, 15,000 and 20,000 updates.

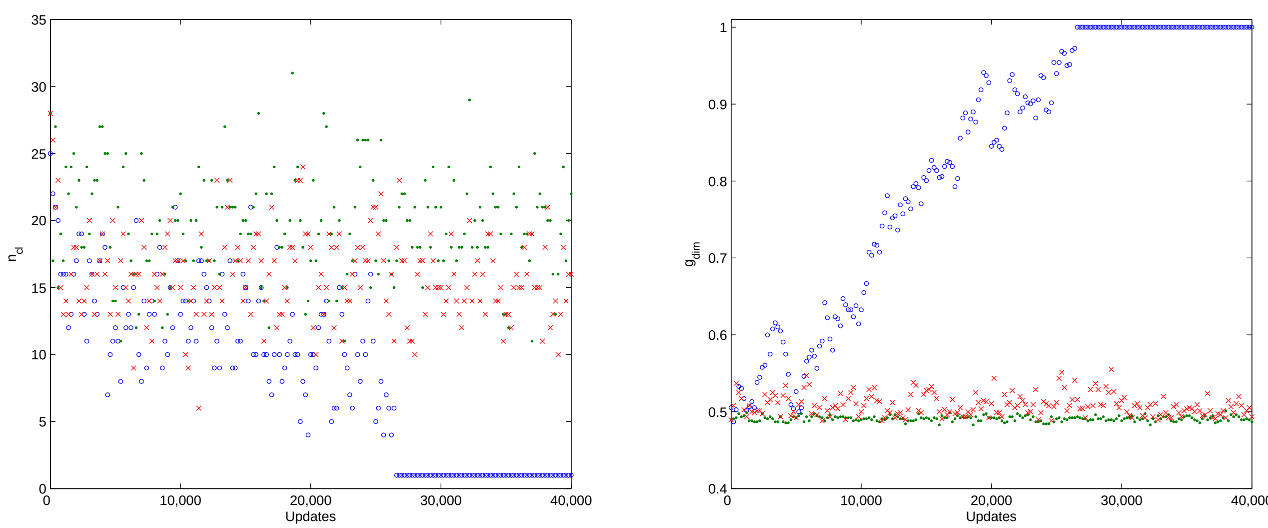

Figure 13: [Left] The number of opinion clusters, n_{cl}, and [Right] the relative giant cluster size, g_{dim}, of a run the voter model (blue circles), FM(i) (green dots) and FM(ii) (red crosses) on the Caltech network.

on the lattice. We define an opinion cluster on a network to be a group of nodes such that all nodes in the group have the same spin and for any two nodes in the group there exists a path between them only passing through nodes within the group. Another measure we use as an indicator of the dynamics is g_{dim}, the relative giant cluster size, the size of the largest opinion cluster divided by the total number of nodes [19]; Figure 8 (Right) plots this measure for a run of our three models on the lattice.

Figures 7 and 8 show that on the one dimensional lattice for the voter model large clusters quickly emerge, the clusters of one of the two spins spread, merging and proceeding to consensus. For FM(i) there are typically large numbers of clusters, however cluster structure is not preserved by the dynamic and clusters do not persist through time; for FM(ii), on the other hand, the frequency of small clusters is greatly reduced and clusters persist for much longer. FM(i) and FM(ii) were also run on a 30x30 two dimensional lattice. Figures 10-12 provide a visualization in a similar fashion to that we described for one dimensional lattices - a 30x30 grid represents the lattice, each square of the grid corresponding to a node and the adjacent squares (above, below, left and right) to the node's neighbours. The voter model's behaviour on the lattice has been very well studied in statistical physics, it is observed that as the model is run opinion clusters grow but the interfaces are very rough. Figures 9-12 demonstrate that for FM(i) also cluster boundaries are not particularly clear and, as in the 1D case, many small clusters, often of single nodes, persist over time; for FM(ii) there are fewer clusters interfaces and their boundaries are somewhat sharper.

Finally, FM(i) and FM(ii) were run on the Caltech network illustrated in Figures 14 and 15 respectively, where the colour lines plot the proportion of upspin nodes in each of the districts over time, and the black line the global proportion of upspin nodes. The overall proportion of upspin nodes varies very little from 50% for FM(i) whereas for FM(ii) the proportion fluctuates more freely around this mark. Figure 13 demonstrates that both models produce a similar number of opinion clusters, the largest of which constituting just over half of the nodes. In fact when we look at the dynamics more closely we find that there are always two large clusters: one voting Democrat, the other Republican, the remaining opinion clusters being very small, often of just one or two nodes. It is clear from Figures 14 and 15 that in both cases the vote within each community can vary from the proportions of the overall vote, and this is the effect we wished to model - the popular vote will return a winner for one of the two parties, but within some of the communities the other party will in fact be the majority. If we compare the dynamics of the vote within our districts when FM(i) and FM(ii) are run to the vote in our earlier artificial

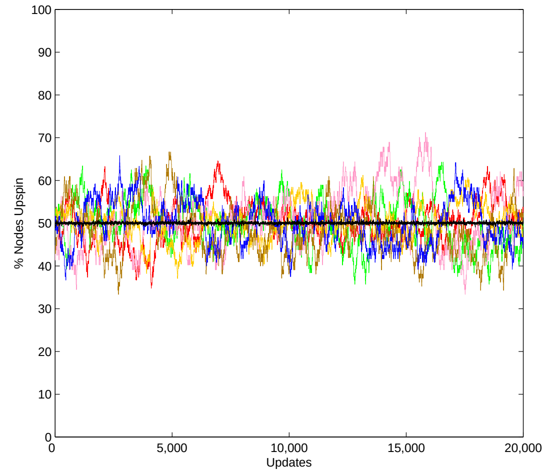

Figure 14: [Colour] A run of FM(i) on the Caltech network. Each coloured line expresses the proportion of upspin nodes (or votes for the Democrats) in one of the communities, the black line represents this figure for the whole network. A line showing the proportions of the vote in the two node District 4 is omitted for clarity.

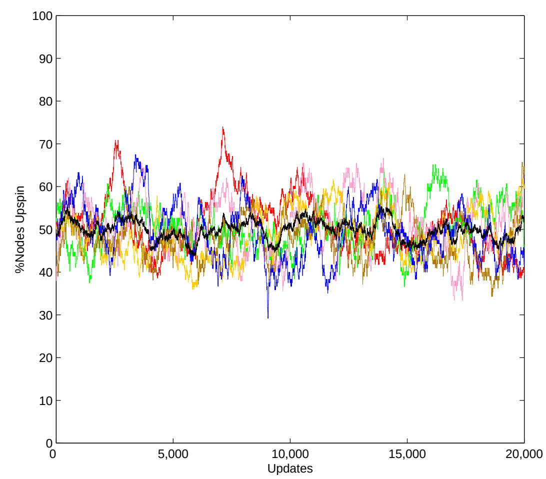

Figure 15: [Colour] A run of FM(ii). Each coloured line expresses the proportion of upspin nodes (or votes for the Democrats) in one of the communities, the black line represents this figure for the whole network. A line showing the proportions of the vote in the two node District 4 is omitted for clarity.

communities when the original voter model was run (see Figure 6), we note that the vote across communities varies much greater in the latter situation. So although the addition of our feedback mechanisms to the voter model is able to simulate local variations from a global vote of around 50% for each party, if we have a very strong community structure to begin with the original voter model alone is successful at producing concentrations of support for the two parties within different communities.

CONCLUSIONS & FURTHER WORK

Using the Facebook networks of three American colleges we modelled two-party elections and the phenomenon that people belonging to the same sphere of influence tend to hold similar political opinions. Using Newman's eigenvector method to detect community structure in our networks, and the voter model as our primary opinion dynamic we found that one requires relatively strong community structure to observe concentrations of different opinions within different communities. By the addition of random flips to the voter model via feedback mechanisms we were able to simulate communities as electoral districts, keeping the global vote fairly similar for each party whilst producing different results in different districts.

Further work should include the investigation of the voting patterns on other networks. For instance we wish to model on larger networks, which were not practically possible given our little time and computing power, as well as on benchmark networks; in this manner we may be able to establish more fully the link between local modularity and fluctuations in voting proportions. Furthermore one could investigate other binary variable opinion dynamics such as the Sznadj model, or other feedback mechanisms.

ACKNOWLEDGEMENTS

I would like to extend my thanks to the EPSRC for their pecuniary support which made my research possible; to Nick Jones and Mason Porter for their mathematical and personal support which made my research both productive and enjoyable; and to Mandi Traud for her code for the ringplots as well as her advice on how to use it.

References

[1] Strogatz S.H. Exploring complex networks. *Nature*, 410(6825):268-276, 2001.

[2] Mark E.J. Newman. The structure and function of complex networks. *SIAM Review*, 45(2):167-256, 2003.

[3] Claudio Castellano, Santo Fortunato, Vittorio Loreto. Statistical physics of social dynamics. *to appear in Reviews of Modern Physics, arXiv: 0710.3256*, 2007

[4] Marta Gonzalez. *Contact Networks of Mobile Agents and Spreading Dynamics*, Ed. Logos-Verlag, 2007.

[5] Chung-I Chou, Sai-Ping Li. Taming the Gerrymander - Statistical physics approach to Political Districting Problem. *Physica A: Statistical Mechanics and its Applications*, 369(2):799-808, 2006.

[6] Anuj Mehrotra, Ellis L. Johnson, George L. Nemhauser. An Optimization Based Heuristic for Political Districting. *Management Science*, 44(8):1100-1114, 1998.

[7] J.A. George, B.W. Lamar, C.A. Wallace. Determining New Zealand electoral districts using a network-based model. *Proceedings of the Operational Research Society of New Zealand*, 29:276-283, 1993.

[8] Micah Altman. Modelling the Effect of Mandatory District Compactness on Partisan Gerrymanders. *Political Geography* 17(8):989-1012, 1998.

[9] Michel Balinski. Fair Majority Voting (or How to Eliminate Gerrymandering). *The Mathematical Association of America*, 115:97-113, 2008.

[10] Michel Balinski, H.P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*, Yale University Press, New Haven, CT, 1982; 2nd ed., Brookings Institution Press, Washington D.C., 2001.

[11] Brian Hayes. Machine Politics. *American Scientist*, 84(6):522-526, 1996.

[12] Micah Altman, Karin Mac Donald, and Michael P. McDonald. Pushbutton Gerrymanders: How computers changed redistricting, in *Party Lines: Competition, Partisanship, and Congressional Redistricting*, Thomas E. Mann and Bruce E. Cain (ed.s), Brookings Press, 2005.

[13] Amanda L. Traud, Eric D. Kelsic, Peter J. Mucha, Mason A. Porter. Community Structure in Online Collegiate Social Networks. *unpublished, arXiv:0809.0690v1*, 2008.

[14] V. Sood, S. Redner. Voter Model on Heterogeneous Graphs. *Phys. Rev. Lett.*, 94(178701), 2005.

[15] Santo Fortunato, Claudio Castellano. Community Structure in Graphs. Encyclopedia of Complexity and System Science, B. Meyers, ed., Springer Verlag, Heidelberg, Germany, 2008 (arXiv:0712.2716).

[16] Mark E.J. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(036104), 2006.

[17] Aaron Clauset. Finding local community structure in networks. *Physical Review E*, 72(026132), 2005.

[18] B. Granovsky, N. Madras. The noisy voter model. *Stochastic Processes and their Applications*, 55:23-43, 1995.

[19] Floriana Gargiulo, Alberto Mazzoni. Can extremism guarantee pluralism? *unpublished, arXiv:0803.3879v3*, 2008