Legislatures as spin glasses

Abstract

I present a method for partitioning legislatures into subsets of legislators that vote in a similar fashion by modelling legislatures as Potts spin glasses. I apply the method to find partitions of the UK House of Commons and the US Senate. I explain how the method is related to methods designed to find partitions of networks into densely connected subsets of nodes.

1 Introduction

Particles possessing a magnetic moment are often known as spins [1]. A spin may interact with other spins, and depending on the interaction energy, the interaction may be ferromagnetic, in which case spins attempt to align, or antiferromagnetic, in which case spins attempt to anti-align. A spin glass is a system of interacting spins that receive conflicting relative ordering instructions due to the presence of both ferromagnetic and antiferromagnetic interactions [2]. The term 'spin glass' is derived from the fact that below a freezing temperature, spin glasses exhibit non-periodic freezing of spins into clusters reminiscent of the amorphous freezing of atoms in a conventional glass [1]. A commonly studied spin glass is $\mathrm{Cu}_{1-x}\mathrm{Mn}_x$ with $x \ll 1$, where the magnetic Mn ions (the spins) are distributed randomly throughout the non-magnetic Cu lattice. The interaction energy between Mn ions is a function of the distance of separation, so that some pairs of ions interact ferromagnetically, whilst others interact antiferromagnetically [10].

Spin glasses exhibit several interesting behaviours: frustration, the inability to minimize the interaction energies between all spins simultaneously; slow dynamics at low temperature due to the development of a complicated 'energy landscape'; and a sharp thermodynamic phase transition despite the mixed interactions [2]. There has been much interest in theoretical models that capture one or more of these behaviours [3, 4]. Several other types of systems featuring interacting agents have been modelled as spin glasses in order to take advantage of theoretical techniques developed in the study of spin glasses, including economic and social systems [5, 6], prebiotic evolution [7], and neural networks [8, 9].

A legislature is a body made up of legislators that vote on bills. Examples include the UK House of Commons (where the legislators are Members of Parliament or MPs) and the US Senate (where the legislators are senators). It is of some interest to find partitions of legislatures into subsets of legislators such that members of a subset vote in a similar fashion. The partitions might highlight legislators who are voting more similarly to members of another party than their own [13], or be used to coarse-grain the legislature to simplify political analysis [12]. Formally, a partition P of a set X is a set of nonempty subsets of X such that every element in X is in exactly one subset [14]. A partition of the set $\{1, 2, 3\}$ is $\{\{1\}, \{2, 3\}\}$, with the other possible partitions being $\{\{1\}, \{2\}, \{3\}\}$, $\{\{1, 2\}, \{3\}\}$, $\{\{1, 3\}, \{2\}\}$ and $\{\{1, 2, 3\}\}$.

The problem of partitioning legislatures has recently been tackled by modelling legislatures as networks. A network consists of a set of nodes (or vertices) connected by edges [15]. In [13], the authors constructed networks from legislation cosponsorship data in the US Congress by letting nodes represent legislators and setting the edge weights equal to the number of cosponsored bills. They proceeded to recursively partition the network into densely connected subsets of nodes, or communities. In both the Senate and House of Representatives, they found that the initial partition into two communities placed Republicans in one community and Democrats in another, with the exception of a few known moderate legislators.

Recently it was pointed out that several community detection methods (methods designed to find partitions of networks into densely connected subsets of nodes) solve a problem equivalent to finding the ground state of a particular type of spin glass [16]. The partitions of nodes into communities correspond to partitions of spins into clusters.

Inspired by this identification and the success of the network partitioning scheme in the US Congress [13], I explain a method for partitioning legislatures into subsets of legislators that vote in a similar fashion by modelling legislatures as spin glasses. In the model, spins represent legislators, the interaction energies between spins are determined by the similarity of the voting records of legislators, and the partitions of the legislature correspond to the partitions of spins into clusters in the ground state of the spin glass.

The paper proceeds as follows. In section 2, I explain the method in more detail and discuss its relationship to community detection methods. In section 3 I demonstrate the application of the method to some model systems. In section 4 the method is applied to find partitions of two real legislatures, the UK House of Commons and US Senate, and to discover factions within a political party. In section 5 I offer some concluding remarks.

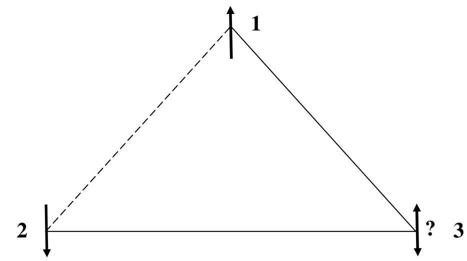

Figure 1: Frustration in a Potts spin glass. Arrows represent spins with the direction of an arrowhead indicating the spin state. The dotted line represents an antiferromagnetic interaction whilst the solid lines represent ferromagnetic interactions. In an attempt to minimize the energy of the system, spins 1 and 2 take different spin states, leaving spin 3 unable to minimize both its interaction energies. No partition of spins into spin states exists such that all interaction energies are minimized simultaneously so the system exhibits frustration.

2 The method

2.1 Potts spin glasses

Specifically, we shall model legislatures as Potts spin glasses. In a q-state Potts spin glass, spins exist in any one of q spin states (point in any one of q directions) such that the interaction energy between spins i and j is $-J_{ij}$ if the spins are in the same spin state (point in the same direction) and zero otherwise [16]. We refer to the quantity J_{ij} as the coupling between spins i and j and the matrix J with elements J_{ij} as the coupling matrix. The Hamiltonian (or energy) of the system is the sum over all interaction energies, and is a function of the partition P of spins into (a maximum of q) spin states. Assigning spins in the i^{th} spin state a quantity σ_i , the Hamiltonian can be written

$$\mathcal{H}(P) = -\sum_{ij} J_{ij}\delta(\sigma_i, \sigma_j) \tag{1}$$

where $\delta(\sigma_i, \sigma_j) = 1$ if $\sigma_i = \sigma_j$ and zero otherwise. If $J_{ij} > 0$ we say there is a ferromagnetic interaction between spins i and j, and the interaction energy is minimized by the spins existing in the same spin state (aligning). If $J_{ij} < 0$ we say there is an antiferromagnetic interaction between spins i and j, and the interaction energy is minimized by the spins existing in different spin states (pointing in different directions). If $J_{ij} = 0$, we say that spins i and j are non-interacting, and the interaction energy is the same whether the spins exist in the same spin state or different spin states.

In general, no partition P exists such that all interaction energies are minimized simultaneously i.e. a Potts spin glass exhibits frustration (figure 1) [22]. Nevertheless, there will be a partition of least energy (the ground state). We refer to the spin states (sets of spins with the same σ_i) in the minimum energy partition of a Potts spin glass as clusters. Spins interacting ferromagnetically have an incentive to belong to the same cluster, whilst spins interacting antiferromagnetically have an incentive to belong to different clusters.

2.2 Legislatures as Potts spin glasses

The method for partitioning legislatures into subsets of legislators that vote in a similar fashion may be outlined as follows:

Imagine the n legislators of a legislature as spins in a n-state Potts spin glass[1]. Introduce ferromagnetic interactions between pairs of legislators that vote similarly, so that they have an incentive to belong to the same cluster, and introduce antiferromagnetic interactions between pairs of legislators that vote dissimi-

[1]We set $q = n$ so that it is possible for each legislators to belong to a cluster on their own.

larly, so that they have an incentive to belong to different clusters. Find the minimum energy partition of legislators into clusters.

We now seek to formalise the method, which we sometimes refer to as the spin glass partitioning (SGP) method. We assume that we have voting data of a legislature in the form of a voting matrix V, where

$$V_{ij} = \begin{cases} 1 & \text{if legislator } i \text{ voted yes on bill } j, \\ -1 & \text{if legislator } i \text{ voted no on bill } j, \\ 0 & \text{if legislator } i \text{ did not vote on } j. \end{cases} \quad (2)$$

From this we wish construct a similarity matrix S such that the element S_{ij} is a measure of the similarity of the voting records of legislators i and j. We define

$$b_{ij}^{\text{same}} = \sum_k \delta(V_{ik}, V_{jk})\delta(|V_{ik}|, 1)\delta(|V_{jk}|, 1) \quad (3)$$

which is the number of bills on which legislators i and j both voted and voted the same way, and

$$b_{ij} = \sum_k \delta(|V_{ik}|, |V_{jk}|)\delta(|V_{ik}|, 1)\delta(|V_{jk}|, 1) \quad (4)$$

which is the number of bills on which legislators i and j both voted. We then let

$$S_{ij} = \begin{cases} b_{ij}^{\text{same}}/b_{ij} & \text{if } b_{ij} > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

The matrix S is symmetric with ones on the diagonal and elements in the range $0 \leq S_{ij} \leq 1$. We define the elements of the coupling matrix J of the legislature as

$$J_{ij} = S_{ij} - \lambda, \quad (6)$$

where we have introduced a parameter λ so that $J(\lambda)$. Pairs of legislators with $S_{ij} > \lambda$ interact ferromagnetically and have an incentive to belong to the same cluster, whilst pairs of legislators with $S_{ij} < \lambda$ interact antiferromagnetically and have an incentive to belong to different clusters. Stated in words, pairs of legislators that voted the same way on more (less) than a fraction λ of the number of bills on which they both voted interact ferromagnetically (antiferromagnetically).

2.3 Connection to community detection methods

We can view the matrix S as an adjacency matrix, or a matrix that represents a network. The element A_{ij} of an adjacency matrix A is the weight of the edge running from node i to node j. If the network is undirected, $A_{ij} = A_{ji}$, and if the network has no self edges, $A_{ii} = 0$. The matrix S therefore represents an undirected network with self edges of unit weight, with nodes corresponding to legislators.

A popular approach to partitioning a network into densely connected subsets of nodes, or communities, is to look for the partition that maximizes a quality function called modularity [21], defined as the sum over all subsets of the difference between the actual sum of edge weights between nodes in a subset minus the expected sum of edge weights between nodes in a subset, divided by sum of all edge weights:

$$Q(P) = \frac{1}{W} \sum_i w_i - [w]_i \quad (7)$$

Here P denotes a partition of nodes into subsets, w_i is the actual sum of edge weights between nodes in the i^{th} subset, $[w]_i$ is the expected sum of edge weights between nodes in the i^{th} subset and the sum runs over all subsets. The variable W is the sum of all edge weights in the network and plays no part in the maximization of Q as it is the same for all partitions. We define $[w]_i$ by picking a null model against which to compare our network. A subset will contribute positively to the modularity only if the sum of edge weights between nodes in the actual network is greater than expected from the null model.

Assigning nodes in the i^{th} subset a community index c_i and defining g_{c_i} as the set of nodes with the same c_i, the modularity can be rewritten as

$$Q(P) = \frac{1}{2W} \sum_{ij} (A_{ij} - p_{ij})\delta(c_i, c_j) \quad (8)$$

where now the sum runs over all nodes in the network and we have assumed that the network has no self edges so that $w_k = \sum_{ij \in g_{c_k}} A_{ij}/2$, $[w]_k = \sum_{ij \in g_{c_k}} p_{ij}/2$ and $W = \sum_{ij} A_{ij}/2$. The term p_{ij} is, by definition, the expected weight of the edge running from node i to node j. The most common choice of p_{ij} (and the one used to detect communities in networks constructed from legislation cosponsorship data in the US Congress [13]) is $p_{ij} = k_i k_j / 2W$, where $k_i = \sum_j A_{ij}$ is the degree of node i. This choice of p_{ij} corresponds to a null model in which the edge weight of the network is distributed randomly subject to the constraint that the degrees of nodes in the null model are the same as those in the network [21]. Several other null models have been considered [16, 17, 19, 20].

Comparing equations (1) and (8), we see that the modularity Q is related to the Hamiltonian of a Potts spin glass $\mathcal{H}$ by

$$\mathcal{H}(P) = -2WQ(P) \quad (9)$$

with the choice $\sigma_i = c_i$ and

$$J_{ij} = A_{ij} - p_{ij}. \quad (10)$$

This connection, pointed out in [16], means that the partition of nodes into communities that maximizes the modularity is the same as the partition of spins into spin states that minimizes the energy of a Potts spin glass with the couplings (10). The coupling between nodes is ferromagnetic where the weight of an edge is greater than expected and antiferromagnetic where the weight of an edge is less than expected.

Recalling our choice of couplings (6), we see that the method we have described for partitioning legislatures is equivalent to detecting communities by modularity maximization in a network described by the matrix S with the choice $p_{ij} = \lambda$. This choice of p_{ij} corresponds to a non-standard null model in which every node is connected to every other node by an edge of weight λ.

2.4 Finding the minimum energy partition

If we are to apply our method, we need a way of finding the minimum energy partition of the Hamiltonian (1) with the couplings (6).

We consider the minimum energy partition in different domains of λ. For $\lambda < 0$, all interactions are ferromagnetic and the minimum energy partition consists of a single cluster of n legislators, whilst for $\lambda > 1$, all interactions are antiferromagnetic and the minimum energy partition consists of n clusters each with a single legislator. For $0 \le \lambda \le 1$, the minimum energy partition depends on the precise form of $J(\lambda)$, and to be sure we have found the ground state we must trial every possible partition of spins into clusters [2] [23].

The number of ways of partitioning a set of n elements into nonempty subsets is given by the Bell number $B_n \approx 10^{13}$ for $n = 20$ [24]. The Bell numbers increase at least exponentially in n which means finding the minimum energy partition by an exhaustive search over all possible partitions takes at least an exponential amount of time [23].

[2] This is not quite true: we can be sure that spins interacting only antiferromagnetically should end up clusters on their own, whilst spins interacting with only one other spin by means of a ferromagnetic interaction should end up in the cluster of that spin. There is no need to trial partitions not satisfying these conditions.

The problem is one of a class of problems in combinatorial optimization, or problems where we are looking for the best possible solution from a finite set of feasible solutions. Various approximate methods, or heuristics, have been developed that sacrifice the guarantee of finding the optimal solution for the sake of getting good solutions in a significantly reduced amount of time [25]. Several of these have been implemented in the context of modularity maximization, including spectral bisection [21], greedy optimization [23, 36], tabu search [17], simulated annealing [16] and genetic [35] algorithms. In this report we look for ground states using the following greedy algorithm[3], which can be viewed as a generalization of [36] to any form of couplings.

The algorithm relies on an expression for the energy change of moving a spin z from a subset X to a subset Y that can evaluated quickly. Before the move, the energy of the system is

$$\mathcal{H}_{\text{before}} = -\sum_{ij\in X} J_{ij} - \sum_{ij\in Y} J_{ij}, \quad (11)$$

where we have neglected the contributions from subsets other than X and Y. After the move, the energy is

$$\mathcal{H}_{\text{after}} = -\sum_{ij\in X\setminus\{z\}} J_{ij} - \sum_{ij\in Y\cup\{z\}} J_{ij}, \quad (12)$$

where $X \setminus \{z\}$ is the set difference of X and z and $Y \cup \{z\}$ is the union of Y and $\{z\}$ [14]. We have

$$\sum_{ij\in X\setminus\{z\}} J_{ij} = \sum_{ij\in X} J_{ij} - \sum_{i\in X} J_{iz} - \sum_{j\in X} J_{zj} + J_{zz} \quad (13)$$

and

$$\sum_{ij\in Y\cup\{z\}} J_{ij} = \sum_{ij\in Y} J_{ij} + \sum_{i\in Y} J_{iz} + \sum_{j\in Y} J_{zj} + J_{zz} \quad (14)$$

so that the change in energy $\Delta\mathcal{H} = \mathcal{H}_{\text{after}} - \mathcal{H}_{\text{before}}$ from moving spin z can be written

$$\begin{aligned}\Delta\mathcal{H} = &\sum_{i\in X} J_{iz} + \sum_{j\in X} J_{zj} \\ &- \sum_{i\in Y} J_{iz} - \sum_{j\in Y} J_{zj} - 2J_{zz}.\end{aligned} \quad (15)$$

This expression involves no set algebra (no unions nor differences) and involves sums over just two subsets X and Y. It can thus be evaluated extremely quickly.

The algorithm (coded in MATLAB) starts by placing each spin in a spin state on its own. The energy

[3] A greedy algorithm is an algorithm that always takes the best immediate, or local, solution while finding an answer [37].

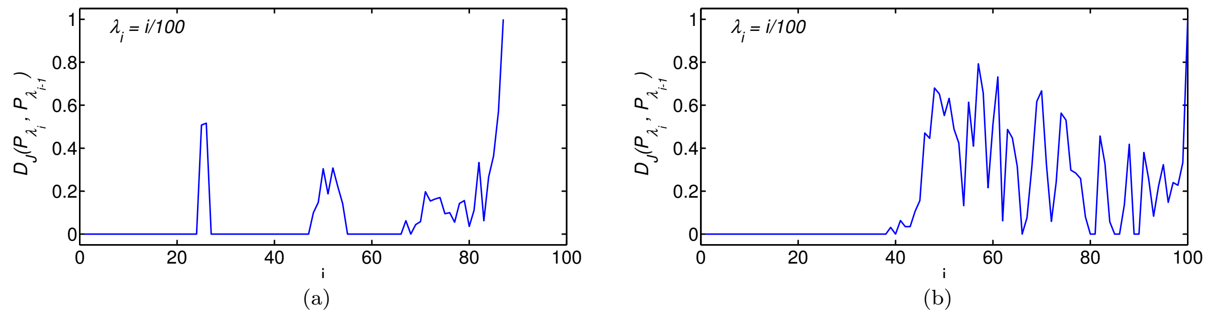

Figure 2: Plot of the Jaccard distance D_J between successive minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} of (a) the artificial legislature featuring parties and coalitions described by (18) (b) the artificial legislature with no political organization described by (19).

of this partition is $-\sum_i J_{ii}$. We consider the energy change (15) that would result from placing spin i in a subset with spin j for all spins j satisfying $J_{ij} > 0$. The spin i is moved to the subset for which the decrease in energy is maximized, unless no move can result in a decrease in energy in which case spin i is not moved. This process is applied repeatedly and sequentially as in [36] until no moves of a single spin can result in a decrease in energy.

Labelling the partition at this point P, we construct a new coupling matrix J' with elements

$$J'_{ij} = \sum_{l \in X_i, m \in X_j} J_{lm} \tag{16}$$

where X_i is the i^{th} subset of P. By applying the first part of the algorithm to this coupling matrix, we determine whether the energy can further be reduced by merging entire subsets of spins. This process is also iterated. The algorithm stops when no merging of subsets can reduce the energy. In subsequent sections, when reference is made to the 'minimum energy partition', it is understood that this refers to the partition of least energy of the partitions returned by the greedy algorithm described above and the spectral bisection algorithm introduced in [21] (also coded in MATLAB).

3 Artificial legislatures

In equation (6) we defined the coupling matrix $J(\lambda)$ up to a constant λ. There is no single proper value of λ, and to take full advantage of the method, one should find minimum energy partitions throughout the range $0 \leq \lambda \leq 1$. Outside this range, the minimum energy partitions are known (and uninteresting) as discussed in section 2.4. Minimum energy partitions that persist over a range of λ are of particular interest, as they correspond to partitions of the legislature that do not break up as the incentives for legislators to belong to the same cluster decrease and the incentives for legislators to belong to different clusters increase. Several authors have proposed similar techniques for finding partitions of networks corresponding to different hierarchical levels [16, 17, 18].

To pick out persistent minimum partitions, we need some way of comparing the similarity of different partitions of the same set. Several indices have been designed for this purpose [26]. In this paper we use the Jaccard distance $D_J(A, B)$, defined as

$$D_J(A, B) = \frac{r_{A!B} + r_{B!A}}{r_{A!B} + r_{B!A} + r_{A\&B}} \tag{17}$$

where $r_{A!B}$ is the number of pairs of elements in the same subset in the partition A and in different subsets in the partition B and $r_{A\&B}$ the number of pairs of elements in the same subset in A and the same subset in B. Smaller values of D_J indicate greater similarity of the partitions A and B, with $D_J(A, B) = 0$ when $A = B$.

We demonstrate the application of the method on two 'artificial' legislatures where we directly define the elements of the similarity matrix to simulate different types of political organization. The first artificial legislature consists of 64 legislators $L_1 \ldots L_{64}$, sixteen political parties $R_1 \ldots R_{16}$ and four coalitions $C_1 \ldots C_4$. Each political party is made up of four legislators, $R_i = \{L_{4i-3} \ldots L_{4i}\}$, and each coalition consists of the union of four political parties, $C_i = \{R_{4i-3} \cup \ldots R_{4i}\}$. The elements of the similarity matrix S_{ij} for $i \geq j$ are defined as

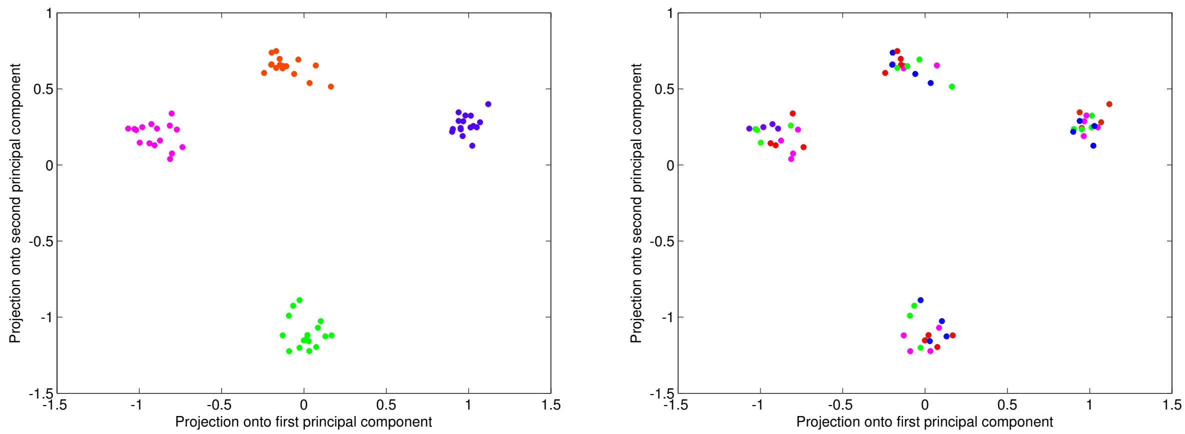

(a) (b)

Figure 3: Projections of the positions of legislators in the first artificial legislature onto the principal components of the similarity matrix (a) coloured by coalition (b) coloured by party.

$$S_{ij} = \begin{cases} 1 & i = j, \\ 0.75 + r & L_i, L_j \in R_k,\ i \neq j, \\ 0.5 + r & L_i, L_j \in C_k,\ L_i, L_j \notin R_k, \\ 0.25 + r & L_i, L_j \notin C_k. \end{cases} \tag{18}$$

We then set $S_{ij} = S_{ji}$ for $i < j$ so that S is symmetric with ones on the diagonal as is the case for similarity matrices constructed from voting data. Here r is a random number taken from a uniform distribution on the interval $[-0.125, 0.125]$. We shall show that by varying λ one can recover both the partitions $P = \{R1, R2 \ldots R16\}$ and $P = \{C1, C2, C3, C4\}$.

Figure 2a shows a plot of the Jaccard distance D_J between successive minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} for $i = 0, 1, 2 \ldots 100$, $\lambda_i = i/100$. For $\lambda > 0.86$, the minimum energy partition consists of each legislator in a cluster on their own so that there are no pairs of legislators and D_J is undefined.

We observe three regions where $D_J = 0$ so that there are three persistent minimum energy partitions. The transitions between the persistent partitions occur when $\lambda \approx 0.25$ and $\lambda \approx 0.5$, reflecting the construction of the similarity matrix (18). The first persistent partition is a continuation of the partition of legislators into a single cluster that occurs for all $\lambda < 0$. The second persistent partition is $P = \{R1, R2 \ldots R16\}$, the partition of legislators into parties, whilst the third is $P = \{C1, C2, C3, C4\}$, the partition of legislators into coalitions. The partitions in the transitional regions around $\lambda = 0.25$ and $\lambda = 0.5$ consist of a mix of persistent partitions. As λ is increased above $\lambda = 0.75$, most of the interactions between legislators are antiferromagnetic and clusters consist of fewer and fewer legislators until at $\lambda = 0.86$, every legislator is in a cluster on their own.

The second artificial legislature also consists of 64 legislators but this time we suppose that there is no political organization, defining the elements the similarity matrix for $i \geq j$ as

$$S_{ij} = \begin{cases} 1 & i = j, \\ r & \text{otherwise.} \end{cases} \tag{19}$$

As before, $S_{ij} = S_{ji}$ for $i < j$ but now r is a random number taken from a uniform distribution on the interval $[0, 1]$. The plot of $D_J(P_{\lambda_i}, P_{\lambda_{i-1}})$ for $i = 0 \ldots 100$, $\lambda_i = i/100$ is shown in figure 2b. This time we observe no regions where $D_J = 0$ apart from the initial stretch corresponding to a partition into a single cluster.

3.1 Visualising legislatures

We can imagine the rows of an $n \times n$ similarity matrix as defining the coordinates of legislators in a n dimensional space. The coordinate of the i^{th} legislator on the j^{th} axis is S_{ij}, so that legislator i has a coordinate of unity on the i^{th} axis ($S_{ii} = 1$) and all legislators appear in the first (hyper)quadrant ($S_{ij} \geq 0$). It would be useful to have some way of visualizing the n-dimensional system to be able to make predictions about the number and nature of persistent minimum energy partitions.

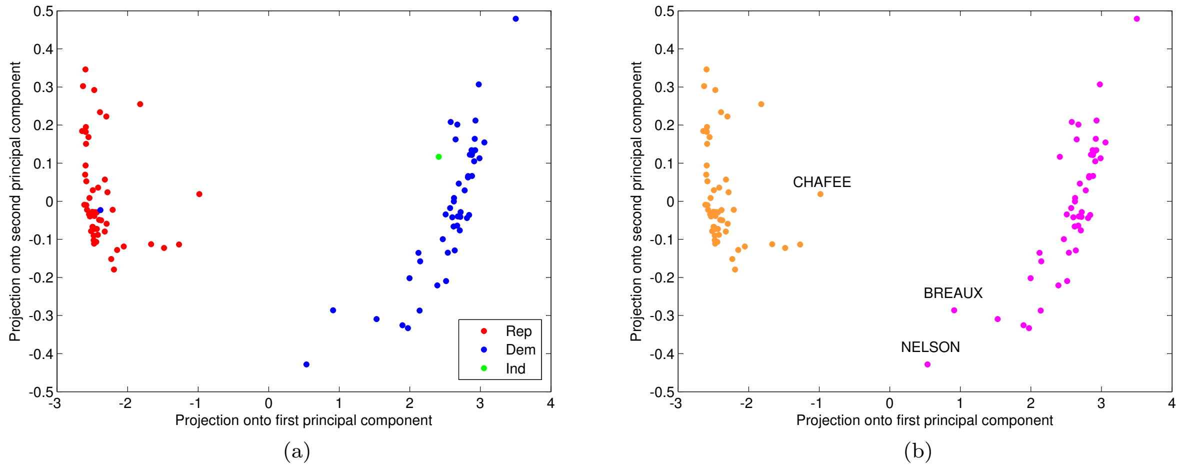

Figure 4: Projections of the positions in legislators in the 108$^{\text{th}}$ US Senate onto the principal components of the similarity matrix (a) coloured by party (b) coloured by the minimum energy partition for $0.42 \leq \lambda \leq 0.72$.

We perform (statistical) dimensional reduction into two dimensions by projecting the coordinates of legislators on to the first two principal components. The principal components are the (orthogonal) leading eigenvectors of the covariance matrix [27]. The first principal component is the direction of maximum variance, and it is the best fit line in the sense that it minimizes the sum of squares of perpendicular distances between the data points (here the positions of legislators) and the line[4]. The second principal component is the direction of maximum variance perpendicular to the first principal component. Projections should be treated with caution, for although legislators that appear close together are close in the directions of greatest variance, they may still be far apart in several other directions.

The projection of the artificial legislature described by the similarity matrix (18) is shown in figure 3a. In this figure, legislators are coloured by coalition and the coalitions are easily distinguishable. However, colouring the legislators by party as in figure 3b, we see that the party structure of the legislature is not at all obvious. This example acts as a warning that projections onto the principal components can hide relevant structure [38].

[4]In fact in least squares fitting it is more common to minimize the sum of squares of *vertical* distances between the data points and the line.

4 Real legislatures

In this section we shall apply the SGP method to partition the US Senate and the UK House of Commons into subsets of legislators that vote in a similar fashion. Each legislator in these legislatures is member of one of a small number of political parties. We would expect legislators of the same party to cluster together for two reasons. Firstly, legislators belonging to the same party tend to share similar political opinions which should translate into similar voting habits. Secondly, members of parties in these legislatures are 'whipped' (told how to vote by the party leadership) [28, 29].

The US Senate is the upper house of the United States Congress. The 108$^{\text{th}}$ Senate (the legislative session between January 2003 and December 2004) featured 100 senators, 51 belonging to the Republican Party (Rep), 48 belonging to the Democratic Party (Dem) and 1 independent (Ind). We base our analysis on a voting matrix obtained from [30] of the recorded votes (roll-call votes) of the 100 senators voting on 675 bills. From this voting matrix we construct a similarity matrix following the process described in section 2.2.

From the fact that almost all legislators belong to one of two political parties and the parties themselves have different political ideals, we might expect a single persistent minimum energy partition approximately equal to {{Rep}, {Dem}}. This assertion is supported the projection of the legislature onto the first two principal components (figure 4a), in which we observe that most legislators are relatively close to

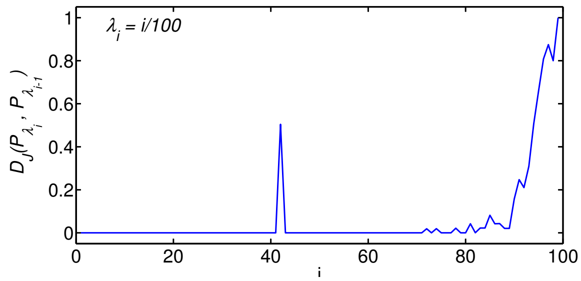

Figure 5: Plot of the Jaccard distance D_J between successive minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} of the 108th US Senate.

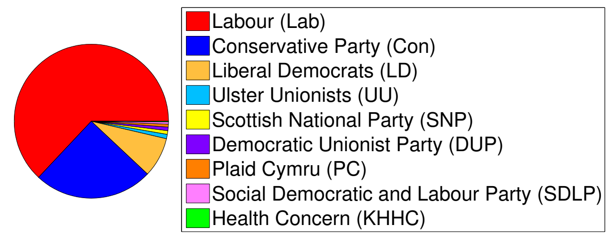

Figure 6: Pie chart showing the proportion of MPs from each party in the UK House of Commons 2001 – 2005. The party abbreviations are used in the text.

members of their own party and relatively far away from members of the other party. A notable exception is the Democrat placed firmly amongst the Republicans: this is Zell Miller, whose anomalous voting behaviour is well known [39].

Figure 5 shows a plot of the Jaccard distance D_J between the minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} for $\lambda_i = i/100$, $i = 0, 1, 2 \ldots 100$. As expected, there is a single persistent minimum energy partition of the legislature corresponding closely to {{Rep}, {Dem}}. This partition, occurring for $0.43 \leq \lambda \leq 0.72$, is shown in figure 4b. As λ is increased above $\lambda = 0.72$, senators begin to break off from the two main clusters. The first three senators to leave the main clusters are labelled in figure 4b. All three were placed in communities predominantly made of up senators from the opposing party in the network study of legislation cosponsorship data [13]. As expected, by $\lambda = 1$ all legislators form clusters on their own.

Next we turn to UK House of Commons, the lower house of the Parliament of the United Kingdom. We consider voting data from the amalgamation of the five parliamentary sessions between June 2001 and April 2005 (Tony Blair's second term as Prime Minister) over which time 657 MPs recorded their votes on 1246 bills. We construct the similarity matrix from a voting matrix of the recorded votes obtained from [31].

Figure 6 lists the abbreviations used to refer to the parties and shows a pie chart illustrating the relative numbers of MPs belonging to each party. The few MPs that switched parties over the period are assigned to the first party of which they were members. We see that in this period the Commons featured three large parties (Lab, Con, LD) in contrast to the two-party system in the 108th Senate. The largest party, Labour, was the governing party. Armed with this knowledge, we may predict the existence of two

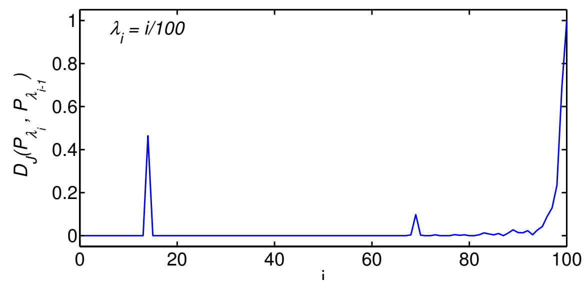

Figure 7: Plot of the Jaccard distance D_J between successive minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} of the UK House of Commons 2001 – 2005.

persistent minimum energy partitions: first, a partition that separates the three major parties, {{Lab}, {Con}, {LD}}; second, a partition {{Lab}, {Con, LD}} corresponding to the Conservatives and Liberal Democrats uniting against the majority Labour government. We have neglected minor parties in these predictions, and it is of interest to see how they fit in to the partitions.

To obtain a further idea of the minimum energy partitions we can expect to find, we study a projection of the positions of legislators onto the first two principal components coloured by party (figure 8a). The projection backs our prediction of a partition approximately equal to {{Lab}, {Con}, {LD}} and is not inconsistent with the prediction of a coarser partition approximately equal to {{Lab}, {Con, LD}}. The location of legislators belonging to minor parties in the projection suggests that in these partitions, the minor parties are likely to be included in clusters along with the major ones rather than forming separate clusters.

Figure 7 shows a plot of the Jaccard distance D_J between successive minimum energy partitions for

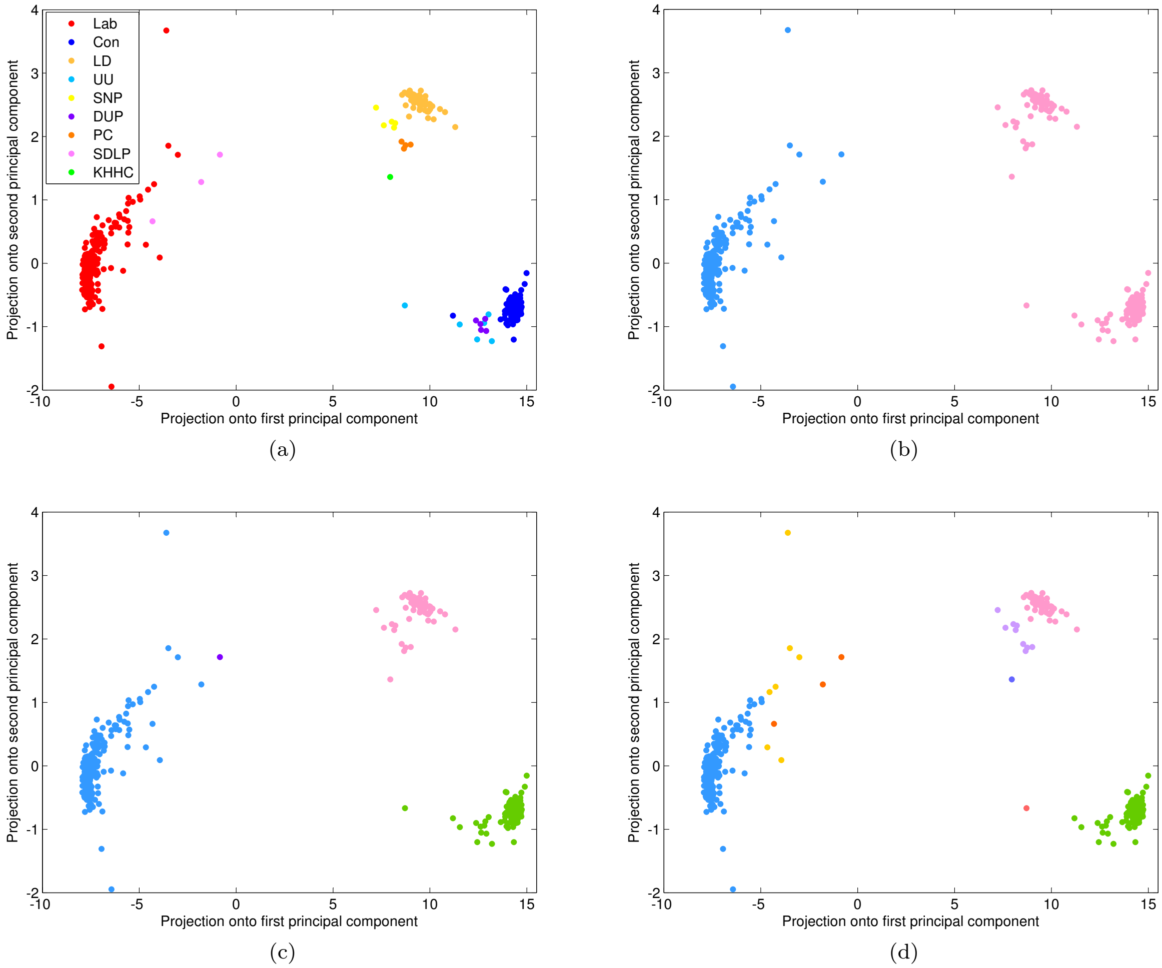

Figure 8: Projections of the positions of legislators in the UK House of Commons 2001 – 2005 onto the principal components of the similarity matrix (a) coloured by party (b) coloured by the minimum energy partition for $0.15 \leq \lambda \leq 0.68$ (c) coloured by the minimum energy partition at $\lambda = 0.70$ (d) coloured by the minimum energy partition at $\lambda = 0.85$.

$\lambda_i = i/100$, $i = 0, 1, 2 \ldots 100$. We are able to pick out two persistent minimum energy partitions. The first of these, for $0.15 \leq \lambda \leq 0.68$, is {{Lab, SDLP}, {other parties}}, shown in figure 8b. This partition matches the prediction of a partition approximately equal to {{Lab}, {Con, LD}}. We may explain the presence of the SDLP MPs in the 'government' cluster by noting that over this period, the SDLP informally accepted the Labour whip (SDLP MPs agreed to vote with Labour) [32].

The second persistent partition is not truly persistent but changes very little across $0.70 \leq \lambda \leq 0.80$. This partition, shown in figure 8c, matches our other prediction and the partition hinted at by figure 8a. We note that the Scottish and Welsh nationalist parties (the SNP and PC) align with the Liberal Democrats whilst the Irish parties (DUP and UU) align with the Conservatives.

The partition for $\lambda = 0.85$ is shown in figure 8d. By this point, the SDLP have separated from Labour, the SNP and PC have formed a cluster on their own and, interestingly, a group of Labour MPs have broken away from the bulk of the party. Upon investigation, this group turns out to consist of known rebel Labour MPs including Jeremy Corbyn, Bob Marshall-Andrews and John McDonnell [33].

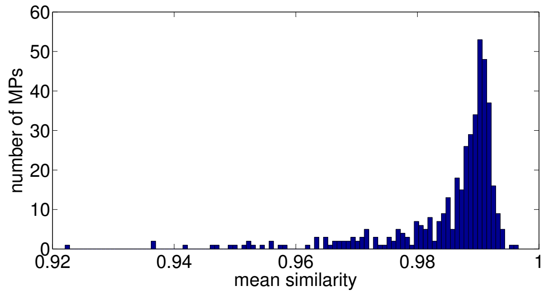

Figure 9: Histogram of the mean similarity $\bar{s}_i$ of members of the UK Labour Party 1997 – 2001. There are 100 bins of equal width between $\bar{s}_i = 0.92$ and $\bar{s}_i = 1$.

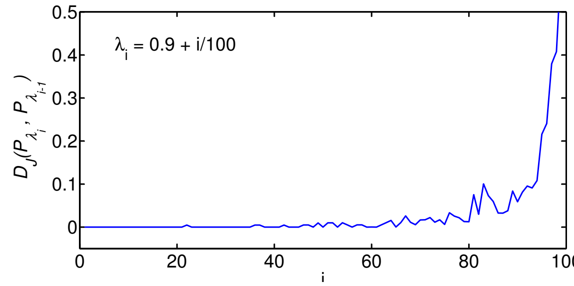

Figure 10: Plot of the Jaccard distance D_J between successive minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i} of the UK Labour Party 1997 – 2001. Note that here $\lambda_i = 0.9 + i/100$.

4.1 Intra-party voting blocs

We now apply our method to a slightly different problem: the identification of intra-party voting blocs, or factions. We attempt to identify factions within the Labour Party in the UK House of Commons during the first Blair government, May 1997 – May 2001.

In a previous study on this topic [12], Quinn and Spirling (Q&S) introduced and applied a nonparametric statistical model based on Commons voting data to find nine sizeable factions of MPs within the party. By a combination of statistical and political analysis, they rated one faction (that included several members of the left-wing Socialist Campaign Group [34]) as by far the most rebellious and a second faction as some deal more rebellious than the others.

We construct a similarity matrix from the recorded votes of 423 Labour MPs on 1279 bills over the period, using data obtained from [31]. Following Q&S, we exclude the two MPs that switched parties and one that never voted. This time, the minimum energy partition consists of a single cluster all the way up to $\lambda \approx 0.9$. This may be explained the distribution of mean similarities of MPs (figure 9), the mean similarity defined as the quantity

$$\bar{s}_i = \frac{1}{n}\sum_{j=1}^{n} S_{ij} \tag{20}$$

The peak is of the distribution is at $\bar{s}_i = 0.99$, indicating that only values of λ close to unity can induce antiferromagnetic interactions between legislators and produce interesting partitions. For this reason, we consider minimum energy partitions P_{λ_i} for $\lambda_i = 0.9 + i/100$, $i = 0, 1, 2 \ldots 100$. Figure 10 shows a plot of the Jaccard distance D_J between the minimum energy partitions $P_{\lambda_{i-1}}$ and P_{λ_i}, where we have restricted the range of the vertical axis to better show the structure of the plot for $\lambda < 0.995$.

In contrast to the plots for the US Senate and the UK House of Commons, there are no sharp peaks separating regions of small D_J that would suggest transitions of significant numbers of MPs in a small range of λ. Rather, the plot suggests more gradual changes in the minimum energy partition as λ is increased.

The first few peaks correspond to MPs breaking off from the main cluster. Interestingly, these MPs leave the main cluster to join the same smaller cluster. By $\lambda = 0.960$, this cluster has increased in size to 16 members, including all but one of the 10 members of the most rebellious faction as identified by Q&S. As λ is increased further, MPs continue to break off from the main cluster, some to form clusters on their own, others to join the first rebel cluster. Around $\lambda = 0.975$ a second rebel clusters forms and grows in size until at $\lambda = 0.986$ it has 33 members, 19 of which feature in the second most rebellious faction of 36 members identified by Q&S. By this point the first rebel cluster has itself broken up, with some members having joined the second rebel cluster and others forming clusters on their own or with a few others. As λ is increased above $\lambda = 0.986$, several sizeable clusters form and break up again, none bearing more than a passing resemblance to the other factions identified by Q&S. By $\lambda = 1$, each MP forms a cluster on his or her own.

5 Conclusion

I have explained a method for partitioning legislatures into subsets of legislators that vote in a similar fashion by modelling legislatures as Potts spin glasses. The method can be used to obtain partitions of interest in real legislatures. I have highlighted the link between the SGP method and modularity maximization in networks and generalized the algorithm [36] to work with any form of couplings.

The SGP method, based on the Potts model, is a hard clustering method (each legislator is either a member of a cluster or not). I speculate that other spin glass models (XY, Heisenberg [11]) might have some application to the fuzzy clustering of legislators, assigning degrees of membership of each legislator to each cluster.

I thank my supervisors, Mason Porter, Jukka-Pekka Onnela and Nick Jones, for many useful discussions. I thank Peter Mucha for the idea of using the Jaccard distance to compare partitions and related MATLAB code.

References

[1] D. Sherrington, arXiv:cond-mat/0512425 (2005)

[2] A. P. Young, Journ. Phys. **95**, 012003 (2008)

[3] S. F. Edwards and P. W. Anderson, Journ. Phys. F **5**, 965 (1975)

[4] D. Sherrington and S. Kirkpatrick, Phys. Rev. Lett. **35**, 26 (1975)

[5] F. Shafee, arXiv.org:physics/0506161 (2005)

[6] P. Contucci and S. Ghirlanda, arXiv.org:physics/0606062 (2006)

[7] K. Nokura, Journ. Phys. A **36** 47 (2003)

[8] D. J. Amit, H. Gutfreund and H. Sompolinsky, Phys. Rev. A **32** 2 (1985)

[9] I. Kanter, Phys. Rev. A **37** 7 (1988)

[10] S. Blundell, Magnetism in Condensed Matter (2001)

[11] J. M. Yeomans, Statisical Mechanics of Phase Transitions (1991)

[12] K. Quinn, A. Spirling, polmeth.wustl.edu/retrieve.php?id=553 (2005)

[13] Y. Zhang, A. J. Friend, A. L. Traud, M. A. Porter, J. H. Fowler and P. J. Mucha, arXiv.org:0708.1191 (2007)

[14] S. Lipschutz, Outline of Set Theory and Related Topics (2005)

[15] M. E. J. Newman, arXiv:cond-mat/0303516v1 (2003)

[16] J. Reichardt and S. Bornholdt, Phys. Rev. E **74**, 016110 (2006)

[17] A. Arenas, A. Fernandez and S. Gomez, arXiv:physics/0703218v2 (2008)

[18] A. Lancichinetti, S. Fortunato and J. Kertesz, arXiv:0802.1218v1 (2008)

[19] P. Ronhovde and Z. Nussinov, arXiv.org:0803.2548 (2008)

[20] T. Heimo, J. Kumpula, K. Kaski and J. Saramaki, arXiv.org:0804.3457 (2008)

[21] M. E. J. Newman, Phys. Rev. E **74**, 036104 (2006)

[22] P. Schiffer, Nature **420**, 35 (2002)

[23] M. E. J. Newman, Phys. Rev. E **69**, 066133 (2004)

[24] E. W. Weisstein, Mathworld: Bell Number (2008)

[25] C. Blum and A. Roli, ACM Comput. Surv. **35**, 3 (2003)

[26] L. Denoeud, H. Garreta and A. Guenoche, Data Science and Classification (2006)

[27] J. Shlens, A Tutorial on Principal Components Analysis (2005)

[28] UK Parliament website, http://www.parliament.uk/

[29] W. J. Oleszek, Journ. Pol. **33**, 4 (1971)

[30] K. T. Poole, http://voteview.com

[31] D. Firth and A. Spirling, http://www.warwick.ac.uk/go/tapir/ (2003)

[32] W. A. Hazleton, Journ. Leg. Stud. **1**, 4 (1995)

[33] P. Cowley and M. Stuart, Dissension amongst the Parliamentary Labour Party, 2001-2005, http://www.revolts.co.uk

[34] Socialist Campaign Group, http://www.poptel.org.uk/scgn/

[35] M. Tasgin, A. Herdagdelen and H. Bingol, arXiv.org:0711.0491 (2007)

[36] V. D. Blondel, J. Guillaume, R. Lambiotte and E. Lefebvre, arXiv.org:0803.0478 (2008)

[37] P. E. Black, Dictionary of Algorithms and Data Structures (2005)

[38] K. Yeung and W. Ruzzo, http://www.cs.washington.edu/homes/kayee/pca/

[39] BBC News, americas/3600710