

Using political science voting models to determine weightings in multi-objective decision problems

Andrea Abel , Salah Sukkarieh

University of Sydney, Australia

a.abel,salah@acfr.usyd.edu.au

Abstract

In this paper we present an approach to the coordination of complex systems operating in a multi-objective domain, which is an extension of the dynamic weighted aggregation method where the weightings are influenced by the decision makers themselves. This is implemented through an information theoretic voting model, taking its inspiration from the political science quantitative models of elections. The example system presented shows this coordination method, where voters vote for the parties (and hence the weightings) corresponding to their preferences, and parties update their policy (weighting) positions learning from the voters.

1 Introduction

Most real world problems are multi-objective, with the objectives often providing sources of conflict.

The decision problem, often referred to as an optimisation problem (from the decision maker's point of view), can be formulated as follows:

$$\max \mathbf{f}(\mathbf{x})$$

where $\mathbf{x} = [x_1, x_2, x_3, \dots, x_n]$ is a vector of n decision variables, and $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), f_3(\mathbf{x}), \dots, f_n(\mathbf{x})]$ a vector of m objective functions, subject to the constraints $g(x) \in X$. Note that optimisation may refer to minimisation or some other definition of optimum.

1.1 Coordination in multi-objective missions

The coordination of entities may be enforced (in a centralised system), through a central 'controller' who receives information from all entities and based on this total information, then delegates commands (perhaps in the form of actions). As the size of the system grows in terms of both objectives and entities, it becomes increasingly more difficult and costly to manage such a

centralised system, hence decentralised approaches to coordination become necessary.

In single and dual objective missions, decentralised approaches to coordination that have been put forward have been primarily based on concepts of bargaining and direct negotiations between entities, see for example [Grocholsky, 2002].

When the number of objectives is extended, together with a significant increase in the number of entities, it may become more efficient to implement behaviour based coordination strategies such as voting, whilst keeping the system decentralised. One avenue which has been extensively explored for decentralised coordination of multiple entities with multiple objectives, is a free market based approach, see the work of Stentz, for example [Zlot, Stentz, 1984].

The purpose of this paper is to present a behaviour based decentralised coordination approach to multi-objective decision problems.

1.2 Approaches to multi-objective decision problems

Since the pioneering work of Rosenthal [Rosenthal, 1984], many methods have been proposed for generating solutions to the multi-objective decision problem, ranging from various types of aggregation methods through to lexicographic ordering, genetic algorithms and voting. For a comprehensive overview of these and other methodologies, see [Coello, 1999].

For our purposes it is of particular interest to look at the weighted aggregation method. The decision problem in this case is stated as:

$$\min/\max \sum_{i=1}^n w_i f(x_i) \quad (1)$$

$$\sum_{i=1}^n w_i = 1 \quad (2)$$

The question becomes; How are the weighting coefficients chosen? In the Constant Weighting Aggregation

method, many arbitrary values for w_i are chosen, and the decision problem solved for each combination. In the Dynamic Weighting Aggregation (DWA) method, w_i is changed continuously, either incrementally (following a sinusoidal pattern) or suddenly (Bang-Bang aggregation method) see [Jin,Olhoffer,Sendhoff, 2001]. Note that in the examples cited in the above, the DWA method accounts for only two objectives and hence two weightings, where $w_2 = 1 - w_1$. Furthermore, the decision makers themselves have no input into the weighting assignment, and the weightings may not necessarily reflect the respective importance of the objective.

1.3 A different approach to dynamic weighting aggregation method

What is required in the dynamic weighting aggregation method, is for the decision makers to be able to influence the weighting assigned to each objective function, and it is proposed that this be done by voting. Thus the weighting is discretely changed at each time step, but not arbitrarily. Additionally, the objective functions are all defined so as to be measured on the same scale allowing the changing weights to reflect the changing importance of each objective.

Since we are considering decision problems with many entities and many (oftentimes competing) objective functions, the novel voting process takes its inspiration from society, which is also characterised by many entities and many objectives. The field of political science provides a rich theory of quantitative modeling of elections and voting, to which we look as a basis for this model. Methods of coordination for robots via voting have been proposed previously, see for example [Pirjanian, 1998], but these have been largely restricted to a single robot with multiple objectives, and have not been modeled using a political science background.

2 Political science approach to modeling elections

Beginning with the works of Downs [Downs, 1957], the field of political science has developed various methods of modeling political elections, the core of which is the fusion of voter behaviour and party behaviour.

2.1 Spatial model of elections

The use of a ‘spatial’ model of elections, based on an m -dimensional Euclidean ‘issues’ space where both voters and parties are assumed to have spatial preferences, has now become widely accepted as a legitimate basis for modeling elections and parliaments [Ordeshook, 1993]. The dimensionality of the ‘issues’ space depends on the number of issues at hand, and in its simplest form it is a uni-dimensional (straight line) model. The spatial

preferences of voters, (termed ‘bliss points’) and the platforms adopted by the parties as well as the parties’ ‘ideal policies’ as represented by policy points, lie within this issue space. Both voter’s bliss points, and party policy points, are a quantitative measure of their stance on a particular issue. Figure 1 gives a visual representation of a spatial model of elections with an issue space of dimensionality three.




Figure 1: Issue space for a three issue election with voter bliss points and party ideal policy points

2.2 Voter model

In its most simplistic form assuming full information, the voter model allows voters to assess each party’s position (and hence vote for its most preferred party), based on a calculation of utility to the voter from voting for that party. Clearly, the utility function must be decreasing in distance from voter bliss point. This utility is most often characterised as the negative of the squared Euclidean distance between voter i ’s bliss point x_{ik} and party j ’s position x_{jk} :

$$U_{ijk} = -(x_{jk} - x_{ik})^2 \quad (3)$$

That is, the utility to voter i of party j on issue k is the distance between party j ’s policy point on issue k and voter i ’s bliss point on issue k , squared.

Another often used form for utility is the ‘city block’ (absolute) distance between the two positions. This is given as:

$$U_{ijk} = -|(x_{jk} - x_{ik})| \quad (4)$$

The total utility for voter i from party j , over all issues is a simple summation of utilities across all issues:

$$U_{ij} = -\sum_k [(x_{jk} - x_{ik})^2] \quad (5)$$

The first extension of this model, as presented in [Gill, 2004], recognises that there may be some non-issue factors affecting voter utility, pertaining perhaps to the individual candidate or party, for example the voter’s perception of a candidate’s honesty. This is modeled as:

$$U_{ij}^{extended} = U_{ij} + O_{ij} \quad (6)$$

where O_{ij} is the term accounting for these non-issue factors.

A further extension of the voter model takes account of uncertainty in the world, thus a voter may only have an expectation about a party platform, meaning only an expected utility can be calculated:

$$E(U_{ij}) = - \sum_k \left[(E(x_{jk} - x_{ik}))^2 - Var(x_{ijk}) \right] + O_{ij} \quad (7)$$

From a political science perspective, this variance term introduces complications, since it cannot be directly observed, and thus an extra modification is made to the expected utility formulation, replacing the variance term with an entropy term H_{jk} [Gill, 2004]:

$$E(U_{ij}) = - \sum_k \left[(E(x_{jk} - x_{ik}))^2 - H_{jk} \right] + O_{ij} \quad (8)$$

where the entropy is calculated based on probabilities p from observations in some questionnaire data set, that is:

$$H_{jk} = -p \log p \quad (9)$$

Note here, that this entropy term is not indexed by i , as it is calculated based on the sample population of the questionnaire, thus it is related to the party in question rather than the individual voter.

The objective for the voter is to vote for the party giving it maximum expected utility, by the above definitions of utility, this occurs where voter bliss point and expected party policy point are coincident on every issue. In this case the utility to the voter is zero (perhaps counter-intuitively), and decreases with any increase in distance.

2.3 Party model

Party models have conventionally been divided into two types, those of the ‘office-seeking’ party, and those of the ‘policy-motivated’ party.

The former has a policy platform (which is a vector of policy points covering all issues), located at the point of maximum expected probability of winning, regardless of party ideology. In a two-party system, the utility accruing to this type of party (as defined by the number of votes by which it is ahead of its rival), is modeled according to [Miller, Stadler, 1998] as:

$$U_j = E_j - E_m \quad (10)$$

$$E_j = \sum_{i=1}^n P(U_{ij} - U_{im}) \quad (11)$$

$$E_m = 1 - E_j \quad (12)$$

where E_j is the expected number of votes for party j and $P(U_{ij} - U_{im})$ is the probability of voting for party j , given the utility difference ($U_{ij} - U_{im}$) between party j and party m .

The policy-motivated party type is more bounded to its ideal platform x_j^{ideal} , and hence utility as defined by [Smirnov, Fowler, 2003] is;

$$U_{jk} = - (x_{win,k} - x_{jk}^{ideal})^2 \quad (13)$$

where $x_{win,k}$ is the winning party policy point on issue k .

Recently, there has been an increase in literature relating to adaptive platform dynamics (see [Stadler, 1998] for an introduction), recognising the ability of parties to implement policy platform changes over time. One approach taken to capture this dynamic process is to consider an office-seeking party, where the party platform changes based on a certain velocity defined as the derivative of the party’s utility function with respect to the party policy point:

$$x'_{jk} = \frac{dU_{jk}}{dx_{jk}} \quad (14)$$

This is referred to as local gradient hill-climbing.

3 Information theoretic model of dynamic spatial elections

The formulation of the dynamic spatial election model in an information theoretic context, adopts a ‘sensor’ based approach to modeling both voters and parties. From the voter’s perspective, the target state for sensing is the platform position of the parties. The party’s target state is however not so intuitive, being its own party platform as it adapts over time. A constant velocity target model is assumed, since we model the target state with a great deal of uncertainty. The target state vector at time t may be defined as:

$$X_j(t) = \begin{bmatrix} x_{j1} \\ x'_{j1} \\ x_{j2} \\ x'_{j2} \\ \dots \\ x_{jk} \\ x'_{jk} \end{bmatrix}$$

where x_{j1} is party j ’s position on issue 1 and x'_{j1} is party j ’s velocity on issue 1.

We examine the information theoretic formulation of voting through the use of a practical example of system comprising five stationary sensors, trying to coordinate over two objectives; target tracking and terrain observation. Thus we have a two issue model, and the above target state is reduced to two issues, x_{j1} and x_{j2} , and their corresponding velocities.

The sensors receive expected utilities from both target tracking and terrain observation in the form of mutual information gain. The objective for each action is to maximise utility, and we define the two objective functions as $f(track)$ and $f(terrain)$. The importance attached to each objective by each sensor is a normalised weighting of the utilities, thus the ‘issue’ space is defined as a weighting $w \in [0, 1]$. The sensor votes for the party whose issue weighting most closely represents theirs.

The overall mission objective function is a weighted aggregation of the individual objective functions, where the weights are dynamically determined at each time step by the voting protocol, indicating the importance of each objective at and also therefore indicating what task each UAV should conduct. The overall objective function can therefore be represented by:

$$\mathbf{f}(\mathbf{x}) = w_{track} \times f(track) + w_{terrain} \times f(terrain) \quad (15)$$

The objective constraints are a one to one sensor to target assignment.

We assume the number of parties is always equal to the number of mission objectives, this allows the ideal platform of each party to capture an ‘extreme’ weighting for each objective. For the sensor system example then, we have a two party model where the ideal platforms of each party are represented by:

$$\begin{bmatrix} x_1^{ideal} \\ x_2^{ideal} \end{bmatrix}_j = \begin{bmatrix} w_{track} \\ w_{terrain} \end{bmatrix}_j = \begin{bmatrix} 0.9 \\ 0.1 \end{bmatrix}_1 = \begin{bmatrix} 0.1 \\ 0.9 \end{bmatrix}_2$$

A party is not a physical object unlike the voters, who are the sensors. Rather, it is an information theoretic algorithm.

Figure 2 captures the essence of this dynamic weighting voting model for the example problem at hand.

3.1 Voter’s understanding of the world

The process model is a standard constant velocity model, thus the target state is defined as:

$$X_j(t+1) = FX_j(t) + w_j(t) \quad (16)$$

where

$$F = \begin{bmatrix} 1 & dt & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & dt \\ 0 & 0 & 0 & 1 \end{bmatrix}$$




Figure 2: The Issue space for sensor multi-objective problem, the two issues are ‘Weight assigned to target tracking’ and ‘Weight assigned to terrain observation’

and $w_j(t)$ is the process noise, which is taken to be an uncorrelated Gaussian sequence with zero mean and variance defined by:

$$Q_j = \begin{bmatrix} 1/3dt^3 & 1/2dt^2 & 0 & 0 \\ 1/2dt^2 & dt & 0 & 0 \\ 0 & 0 & 1/3dt^3 & 1/2dt^2 \\ 0 & 0 & 1/2dt^2 & 1/3dt^3 \end{bmatrix} \sigma_j^2$$

The term σ_j^2 is the target noise variance, which is dependent on the change in party j ’s mandate over the last two time steps, that is;

$$\sigma_j^2 = \frac{votes_j(t-1) - votes_j(t-2)}{votes_j(t-2)} \quad (17)$$

Therefore, σ_j^2 comes from the population of voter’s, and is related to the target (party), rather than to an individual sensor (voter).

The observation model is based on a ‘city block’ distance between party platform and voter bliss point, the observations for voter i of party j being an inverse of this measure (this is because we will be maximising an information space measure of utility based on observation, thus the smaller the city block distance, the larger the observation metric).

$$Z_{ij}(t) = h(x_j(t), x_i(t)) \quad (18)$$

$$Z_{ij}(t) = \left[\frac{1}{1/|(x_{j1}(t) - x_{i1}(t))|} + \frac{1}{1/|(x_{j2}(t) - x_{i2}(t))|} \right] + v_i(t) \quad (19)$$

The term $v_i(t)$ is the observation noise at time t , which is an uncorrelated Gaussian sequence having zero mean, and a variance R dependent on an entropic measure of win likelihood at time t $H_j(t)$, where;

$$H_j(t) = -p_j \log p_j \quad (20)$$

$$p_j = \text{votes}_j(t-1)/\text{vote}_{\text{total}}(t-1) \quad (21)$$

and

$$R_j = \begin{bmatrix} H_j(t) & 0 \\ 0 & H_j(t) \end{bmatrix}$$

$\text{votes}_j(t-1)$ is the number of votes received by party j at time $(t-1)$, and $\text{vote}_{\text{total}}(t-1)$ is the total number of votes at time $(t-1)$. The variance R is thus not dependent on the sensing perception of the voter, but rather on the strength (in terms of votes) of the party being sensed.

The Jacobian of the observation model with respect to the target state is then defined as:

$$\nabla_{x_j} h(x_j(t), x_i(t)) = \begin{bmatrix} -1/(|(x_{j1} - x_{i1})|^2) & 0 \\ 0 & -1/(|(x_{j2} - x_{i2})|^2) \end{bmatrix}$$

Transformation from state space to information space takes its usual form, as do the prediction and update stages of the information filter, as outlined in [Grocholsky, 2002].

The expected utility to the voter is based on the political science definition of utility ('city block distance'), and thus must be directly derived from the observation. It is formulated as:

$$E(U_{ij})(t) = 1/2 \log((2\pi e)^n (\sum y_{ij}(t))) \quad (22)$$

where for a voter interested in estimating four states, $n = 4$, and $y_{ij}(t)$ is the information vector at time t . The voter votes for the party yielding the highest expected utility.

3.2 Party's understanding of the world

Party platforms (as the parties' target state) are also modeled using the standard constant velocity process model. An assumption is made that the party 'type' is a mixture of office-seeking and policy motivated. The utility function for a party j with a given position x_{jk} on issue k is therefore defined as:

$$U_{jk} = -[(x_{win,k} - x_{jk})^2 + (x_{jk}^{\text{ideal}} - x_{jk})^2] \quad (23)$$

where $x_{win,k}$ is the winning party's position on issue k , and x_{jk}^{ideal} is party j 's ideal position on issue k . We must bare in mind that the ideal case (yielding the party maximum utility), is where $x_{win,k} = x_{jk} = x_{jk}^{\text{ideal}}$ for all k . In this situation, the velocity of adaptation also becomes zero.

Initially, the velocity of platform adaptation (the derivative of U_{jk} with respect to x_{jk}), is zero, since ideally the party would like to win based on its ideal platform position. Changes in velocity occur only at the observation stage.

The process model is given in equation 16, with F defined similarly. Q , the variance of target noise changes, in the σ_j term, which now represents the relative change in mandate strength for party j compared to the other party. Thus:

$$\sigma_j^2 = \frac{\frac{\text{votes}_j(t-1) - \text{votes}_j(t-2)}{\text{votes}_j(t-2)}}{\frac{\text{votes}_j(t-1) - \text{votes}_j(t-2)}{\text{votes}_j(t-2)} + \frac{\text{votes}_m(t-1) - \text{votes}_m(t-2)}{\text{votes}_m(t-2)}} \quad (24)$$

The observation model consists of four observations corresponding to the two issues and the two respective velocities. The velocity for adaptation is the derivative of the utility function defined in equation 23. The observation vector for party j is then defined as:

$$Z_j(t) = \begin{bmatrix} |(x_{j1}^{\text{ideal}} - x_{j1}(t)) + (x_{j2}^{\text{ideal}} - x_{j2}(t))| \\ 2(x_{win,1}(t-1) - x_{j1}(t)) + 2(x_{j1}^{\text{ideal}} - x_{j1}(t)) \\ |(x_{win,1}(t-1) - x_{j1}(t)) + (x_{win,2}(t-1) - x_{j2}(t))| \\ 2(x_{win,2}(t-1) - x_{j2}(t)) + 2(x_{j2}^{\text{ideal}} - x_{j2}(t)) \end{bmatrix} + v_j(t) \quad (25)$$

The observation noise, $v_j(t)$ is uncorrelated Gaussian white noise with a variance R dependent on the strength of the previous time step's winning party's platform, $a(t)$. That is;

$$a(t) = -g(t) \log g(t) \quad (26)$$

$$g(t) = \text{vote}_{win}(t-1)/\text{vote}_{\text{total}}(t-1) \quad (27)$$

where vote_{win} is the number of votes received by the winning party of time $(t-1)$, and

$$R_j = \begin{bmatrix} a(t) & 0 & 0 & 0 \\ 0 & a(t) & 0 & 0 \\ 0 & 0 & a(t) & 0 \\ 0 & 0 & 0 & a(t) \end{bmatrix}$$

The corresponding observation Jacobian is:

$$\nabla_{x_j} h(x_j(t), x_j^{\text{ideal}}, x_{win}(t)) = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

3.3 Results

The results presented in the figures below show the dynamic election and voting process. Figure 3 shows the bliss points of each sensor in Issue space over time.

These bliss points correspond to their preferred platform position and is based on the utility they receive from each objective.




Figure 3: The bliss points of each sensor (as determined by their respective weightings of the two objectives; track target, observe terrain) in Issue space

Figure 4 shows the party platforms dynamically adjusting in issue space, as they consider their ideal policy and the winning party’s policy.




Figure 4: Adaptive party platforms; the parties make observations of previous elections, and update their state estimates (weightings for the two objectives) accordingly

We observe that although the two parties initially offered platforms based on their ideal policy points, the ability to update platform positions based on ‘learning’ what the mandate wants, has seen an obvious conver-

gence of policy positions. This policy convergence is a recognised phenomenon in the political science model of two party (office-seeking type) spatial elections. In Downs’ ([Downs, 1957]) classic two party model complete convergence to the median of voter positions is predicted.

On the other hand, for policy motivated parties in two party elections, a certain degree of policy divergence is usually predicted [Wittman, 1990].

It is because of the inherent assumption we have made that the party types are mixed, and furthermore that the utility accruing to a party (and hence velocity) is equally dependent on the distance from the winning platform and distance from ideal platform, that we see only a degree of convergence between party policy platforms.

Figure 5 shows the velocity with which parties change their platforms, as defined by the derivative of their respective utilities.




Figure 5: The velocity with which the parties’ platforms adapt, based on observations from previous election

And finally, Figure 6 shows the tally of votes for each party over time. Given that in this sensor system example party one wins all the time, observing figure 5, and noting the mathematical definition for adaptive velocity, we can conclude that the velocity of adaptation for party one depends only on the difference between its ideal policy and its actual policy. It is this positive velocity that keeps the policy position for party one from diverging away from its ideal policy further than it has.

On both issues, the changes in velocity are initially much larger for party two (since it is affected by both the winning policy and its ideal policy). The change in policy position is thus quite rapid, and the magnitude of the divergence from ideal policy position on both issues is larger than for party one, as is expected.




Figure 6: Vote tally; number of sensors voting for each party at every ‘election’, based on maximising expected utility to the sensor

The implications of this for the dynamic weighting aggregation method are straightforward. The weighting assigned to each objective is given by the platform of party one. In this case, tracking targets has a higher weighting than terrain observation at each time step, and thus optimisation of target tracking takes priority.

4 Conclusion and future work

A dynamic coordination method for multi-objective problems was presented. The formulation (based around a dynamic weighting aggregation approach) uses concepts established in the political science field, namely the modeling of elections, mapped into information space. This information theoretic formulation is then applied to a simplistic multi-objective coordination problem to test its feasibility.

The example allows us to observe the inner working of the dynamic weighting, and demonstrates its usefulness in coordinating a group of sensors within the scope of a two objective mission. The sensors, as voters, do indeed vote for the party closest in weighting to them, whilst the parties adapt their platforms based a combination of ‘learning’ what the mandate wants, and preconceived preferences.

The example presented provides motivation to extend the application of this multi-objective methodology to systems with a large number (in the order of 10-100) of entities trying to coordinate over multiple objectives. It is envisaged that this methodology be applied to large systems of cooperative UAV’s, but is certainly not restricted to this application.

As the system grows in complexity, and more parties are added in accordance with the increased mission objec-

tives, the majoritarian electoral system being used may not yield a clear winning party. In this case, the party model is extended to allow for coalition formation based on the ‘minimal winning connected’ concept, see [Axelrod, 1970].

References

- [Grocholsky, 2002] Ben Grocholsky. Information theoretic control of multiple sensor platforms. PhD thesis, University of Sydney, 2002.
- [Gill, 2004] Jeffrey Gill. An Entropy measure of Uncertainty in Vote Choice Working paper, University of California, Davis, 2004.
- [Miller, Stadler, 1998] John H. Miller, Peter F. Stadler. The dynamics of locally adaptive parties under spatial voting. *Journal of Economic dynamics and control*, 23:171-189, January 1998.
- [Smirnov,Fowler, 2003] Oleg Smirnov, John H. Fowler. Moving with the Mandate: Policy-motivated parties in Dynamic Political Competition. Annual meeting of the Western Political Science Association, 2004.
- [Stadler, 1998] Barbel M.R. Stadler. Adaptive Platform Dynamics in Multi-Party Spatial Voting. *Advances in Complex Systems*,1:1-22,1998.
- [Pirjanian, 1998] Paolo Pirjanian. Multi-objective action selection in behaviour based control. The proceedings of the 6th Symposium for Intelligent Robotic Systems, 1998.
- [Coello, 1999] Carlos A. Coello Coello. A Comprehensive Survey of Evolutionary-Based Multi-objective optimisation techniques . *Knowledge and Information Systems*, 1999.
- [Jin,Olhoffer,Sendhoff, 2001] Yaochu Jin, Markus Olhoffer and Bernhard Sendhoff. Dynamic Weighted Aggregation for Evolutionary Multi-objective Optimisation: Why does it work and how? . Proceeding of the Genetic and Evolutionary Computation Conference, 2001.
- [Downs, 1957] Anthony Downs. An economic theory of democracy . Harper and Row, New York, 1957.
- [Ordeshook, 1993] Peter Ordeshook. The spatial analysis of elections and committees: Four decades of research . Working paper, California Institute of Technology, 1993.
- [Axelrod, 1970] Robert Axelrod. Conflict of Interest. Markham, Chicago, 1970.
- [Rosenthal, 1984] Richard E Rosenthal. Principles of Multi-objective optimization. *Decision Sciences*, 16:133-152.

[Zlot, Stentz, 1984] Robert Zlot, Anthony Stentz. Market-based Multirobot Coordination Using Task Abstraction. International Conference on Field and Service Robotics (FSR), 2003.

[Wittman, 1990] Donald Wittman. Spatial strategies when candidates have policy preferences. In *Advances in the spatial theory of voting*, James Enelow and Melvin Hinich eds., Cambridge University Press, 1990.