

Optimal Voting in Groups with Convergent Interests

James A. R. Marshall

Department of Computer Science / Kroto Research Institute, University of Sheffield
North Campus, Broad Lane, Sheffield S4 7HQ, UNITED KINGDOM
James.Marshall@sheffield.ac.uk

Abstract

Decision-making is crucially important at all levels of biological complexity, from within single-celled organisms, through neural populations within the vertebrate brain, to collections of social organisms such as colonies of ants and honeybees, or societies of humans. What are the prospects for unifying the study of these apparently disparate systems? All can be conceptualised as voting systems at the appropriate level. In this review I will argue that optimality theory can be of fundamental importance in understanding all these systems. In particular I will argue that for groups without conflict of interests, such as neurons and social insect colonies, similar mechanisms could implement statistically optimal decision-making in apparently highly different systems at different levels of biological complexity. I will consider what currency these systems should optimize, and speculate about the possible application of this understanding to the design of voting systems where individual group members' interests are aligned, such as in certain types of human group, and in collectives of robots. I will also consider how established results from economics and political science, notably Arrow's Impossibility Theorem and Condorcet's 'jury theorem', might relate to what we know of social insect voting systems, where interesting effects such as the emergence of collective rationality from the voting of irrational individuals have recently been demonstrated.

Introduction

In considering social decision-making by collections of individuals such as neurons, ants, honeybees, humans or robots, it seems appropriate¹ to quote from Kenneth J. Arrow's Nobel-prize winning paper: '[...] in a generalized sense all methods of social choice are of the type of voting' (Arrow 1950). Arrow's fundamental and general contribution was to show that no fair voting system is possible, under certain 'natural conditions'. This review,

Copyright © 2011, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹ This paper is a contribution to an AAI Spring Symposium hosted at Stanford University, on 'modelling complex adaptive systems as if they were voting processes'; Kenneth J. Arrow's association with Stanford dates back over 60 years to the publication of his seminal paper on voting theory (Arrow, 1950).

however, will proceed as if 'Arrow's impossibility theorem' does not exist, and survey the theory and evidence for optimal collective choice, or voting, in groups with convergent interests, such as populations of neurons or social insects. I shall also consider whether collective decision-making by human and other groups might be informed by this developing theory. I shall return to Arrow's famous result, and discuss it and the aforementioned 'natural conditions' in the context both of both Condorcet's work on voting systems (Condorcet 1785), and collective decision-making.

Optimal Voting by Neural Populations

While perhaps not immediately obvious, decision-making within individuals can involve voting process. Under this view, the individual voters are neurons within neural populations. A well-studied example is the processing of visual stimuli during a motion discrimination task (*e.g.* Shadlen and Newsome 2001). In this task a subject, such as a monkey or human, is presented with a moving-dot field, in which some proportion of the dots are moving coherently in a particular direction, left or right for example, while the remainder move at random. The decision problem is to decide in which direction the majority of the dots are moving, and indicate this choice. Thus, for a binary decision-problem this is a classic two-alternative forced-choice (TAFC) experiment. The difficulty of the decision problem can be varied by adjusting the proportion of the dots that are moving coherently. Additionally, different experimental protocols are possible such as allowing a subject to indicate their choice when they are ready (free-response paradigm) or forcing them to make the decision a pre-specified amount of time after stimulus onset (interrogation paradigm). This motion discrimination problem has become a classic tool for both behavioural and neuroscientific investigations. The neuroscientific approach has shown that two distinct brain regions are involved in these kind of decisions: sensory evidence on the alternative hypotheses (average motion is leftwards *vs.* average motion is rightwards, for

example) is encoded in neural populations within the medial temporal (MT) region of the visual cortex; these populations then excite corresponding populations in the lateral intra-parietal (LIP) region. These LIP populations are decision populations representing accumulated evidence for the different alternatives, that compete with each other to reach a critical decision threshold (Shadlen and Newsome 2001). When one LIP population reaches this decision threshold, the corresponding option is selected. The LIP populations could thus loosely be conceived of as populations of voters for different alternatives, competing to reach the majority threshold required for a decision. In Arrow's terms, this could be considered a 'method of social choice', in which the society is a society of neurons.

As mentioned above, decision-making in the motion-discrimination task has been studied from behavioural as well as neuroscientific perspectives. An influential model of binary decision-making, known as the 'drift-diffusion model' (DDM) (Ratcliffe 1978) can be applied to behavioural data from the motion-discrimination task. This model provides an exceptionally good explanation of reaction-time and choice distributions for experimental subjects. The DDM also has the very interesting property of corresponding to the statistically-optimal decision mechanism for binary-choice problems, the Sequential Probability Ratio Test (SPRT). The SPRT achieves the statistically-optimal compromise between speed and accuracy of decision-making, minimising expected decision-time for any desired expected error rate by using variable decision-thresholds (Wald and Wolfowitz 1946). Recently it has been shown how neurophysiologically-realistic models of the visual cortex decision-circuits described above can be configured to implement the SPRT decision-making strategy (Bogacz *et al.* 2006). This analysis shows that a crucial feature of optimal decision-making in such systems is competition between neural decision populations; mutual inhibition should occur, in which the greater the activation of one decision-population, the greater the inhibitory effect it should exert on competitor populations. Given that this analysis is based on the organisation of the competing voter populations, and largely ignores the details of the individual voters themselves, it seems interesting to consider whether the mechanisms of neural decision-making might capture fundamental principles of 'social choice' that could also be applied to other voting systems.

Optimal Voting in Social Insect Colonies

Social insect colonies, such as those of ants and honeybees, must frequently engage in collective decision-making during their life history. One well-known example of a

collective decision is allocation of colony foraging effort to multiple food sources of varying quality. Once again, this could be considered as a voting process, as individuals recruit others to their food source, using mass-recruitment pheromones in many species of ant, for example. Recently another fascinating collective behaviour has been increasingly investigated, collective nest-choice during colony emigration or fissioning. The decision problem facing the colony is to select the best of the potential nest sites available nearby. In ants and honeybees, individual scouts discover potential sites and perform noisy assessments of their quality. They then seek to recruit others to their preferred site; ants belonging to various species of the genus *Temnothorax* do this using 'tandem-runs' to physically lead another individual to the site (Moglich 1978), while honeybees (*Apis mellifera*) use the famous waggle-dance to communicate a vector to the site to their followers (von Frisch 1967), who can then travel independently to find it. The populations of scouts committed to different alternatives can be thought of as voters, competing with each other to achieve a majority. When a scout population reaches a quorum threshold, implementation of the decision by transportation of the colony to the corresponding alternative is initiated. There is evidence that this quorum can varied to compromise between the speed and accuracy of decision making (reviewed in Visscher 2007). Recent mathematical analysis of models of these collective decision making mechanisms has capitalised on the striking parallels between their organisation, and that of neural models of decision making (see previous section), applying the optimality theory used in neural models to derive predictions for the interaction patterns within insect colonies that should optimise collective decisions (Marshall *et al.* 2009). As for the neural models, the prediction of such optimality analysis is that populations of voters should compete with each other by inhibiting each others' growth; in the social insect case, theory predicts that this should be achieved by scouts committed to one alternative directly attempting to recruit scouts committed to the other (Marshall *et al.* 2009).

Selecting the Appropriate Optimality Criterion

In attempting an optimality analysis of any system, selecting the appropriate optimality criteria is of crucial importance. For most artificial systems, this is relatively straightforward since the intended purpose of the system is well specified. In considering biological or social systems, however, understanding what should be optimised is typically much more subtle. It must be remembered that organisms, for example, should be expected to optimise their fitness. Simple optimality criteria can provide a good

proxy for this, yet care is still needed in identifying those that are appropriate. This is particularly relevant to the present paper. In studying individual-level decision-making, it is typically assumed that it *is* best for individuals to optimise the time taken on average to achieve a required error rate. In the context of single decisions, and the experimental paradigms typically used, in which a reward is only given when a correct choice is made, this seems appropriate. However, animals must have evolved to maximise their fitness based on *sequences* of decisions over their life history. In particular, if an individual is making a decision as to the better of two options, where a fitness reward is received according to the quality of the selected option, error rate can be the incorrect criterion to optimise. Consider the example of such an individual choosing between two alternatives of very similar, or even identical, quality. Then the decision is difficult and as a result takes a long time (if the decision-making mechanism is designed to achieve a certain expected error rate), yet the benefit from making the correct choice is negligible or even non-existent, since both options gave similar or even identical rewards. The time taken over such a decision must then be traded-off against the opportunity cost from not completing the decision quickly and being able to move onto the next decision, or performing some other important behaviour. The appropriate optimality theory is thus more likely to be related to the theory of maximising reward rate in ‘bandit problems’ (see in particular Gittins 1979), than to the SPRT theory.

Optimal Voting by Groups of Robots and Humans

The similarity between the organisational and optimality principles of neural and social insect voting systems has prompted the suggestion that general principles of optimal collective decision-making might exist (Marshall and Franks 2009). What characterises these systems, and enables the application of appropriate optimality theory, is that the individual voters have aligned interests. While individual neurons or insects vote for different alternatives, this is due to uncertainty in the evidence they sample on the qualities of those alternatives, rather than *a priori* differences in preference. For an insect colony, there is an objective difference in nest quality, in terms of the potential for colony survival and reproduction, and it is in the interest of all members of the colony to select the best of the available alternatives. Similarly neurons in the brain have convergent interests, in that making the correct decision on processing a visual stimulus, and taking the appropriate action in response, benefits all of them equally

in terms of the survival and reproduction of the animal they are part of.

It seems, therefore, that other collective decision-making systems whose members have convergent interests, might benefit from being organised in a similar manner to neural populations or social insect colonies. As summarised above, a defining feature of such organisation is that the populations of voters compete with each other to reach a majority, by directly inhibiting each other’s growth. It is tempting to consider whether decision-making in certain kinds of human group, where there is a common goal but individual voters differ in the information they have available to them, might best be organised in a similar fashion to a colony of social insects such as honeybees (see Seeley 2010, for example). Group decision-making by humans is often implemented using committees, in which individuals deliberate collectively and use a rich communication medium, human language, to persuade others that the information they have on the best course of action is more reliable. In marked contrast, the interactions between ant and honeybee scouts seeking to influence each other are much simpler; individuals recruit more strongly the better they think their preferred alternative is, either by performing more waggle-dances, or beginning tandem-running earlier (see Visscher 2007 and references therein). Yet these simple interactions could be sufficient to implement statistically optimal collective decision-making (Marshall *et al.* 2009). This raises the paradoxical possibility that human groups might make more effective collective decisions by engaging in much simpler, less deliberative communication within those groups.

While it may well be impractical to encourage human groups to make better collective decisions by engaging in simpler inter-individual communication, another kind of artificial system may well benefit from a theory of optimal group decision-making. In collective robotics applications, it seems likely that many group behaviours might usefully be organised along the lines outlined above. For example, similar to social insect colonies, groups of robots may need to collectively select a suitable shelter in some circumstances, from several available. Alternatively, in some applications groups may need to achieve consensus on where the entire group should be deployed to undertake some collective task, based only on local assessments by individuals of the need for that task in different locations. Given the typically simple nature of inter-robot communication at present, and their frequently limited computational power, establishing simple patterns of collective interaction may provide an attractive and effective way of optimising group-level behaviour.

Discussion: Optimality vs. Impossibility

In this paper I have briefly reviewed the potential for treating diverse collective decision-making systems using appropriate optimality theory. Links have already been established between decision-making at two levels of biological complexity, the individual and the collective, and in this paper I have proposed that lessons learned from this may also be applicable to artificial systems of human and robot groups. All of these systems can be conceptualised as voting systems, whether the voters be individual neurons, insects, humans or robots. It seems worth considering whether even lower levels of biological complexity might also make use of the same principles; for example, given proposals that single cells can implement Bayesian estimation (Perkins and Swain 2009), might they also employ the principles of optimal ‘collective’ decision-making summarised here, in which case the ‘voters’ involved would presumably be different chemical species, or proteins, or other intra-cellular products involved in a biochemical network. For signal detection by antibody cells, indeed, cross-talk between receptors has been identified as important (reviewed in George, Stark and Cliburn 2005).

As discussed in the introduction, the majority of this paper has proceeded under the assumption that the voting systems being considered *can* be optimal; the different information available to voters can be integrated into a collective decision in an optimal way, whether by minimising decision time for some required error rate, or maximising reward rate over a series of decisions. Yet, as also discussed in the introduction, a classic result in economics appears to show that *no* mechanism for aggregating social preferences can exist (Arrow 1950). How are we to resolve this apparent paradox? The problem is not with Arrow’s formulation of what constitutes a voting system; for example, Arrow’s result requires that individuals have a preference ordering over alternatives, and this is indeed observed in ants (Franks *et al.* 2003). The resolution of the paradox is at the same time both straightforward, and illuminating in the light of recent experimental evidence on collective choice in social insects. In fact, Arrow explicitly excluded groups with convergent interests from the remit of his theorem. As he pointed out in introducing his result (Arrow 1950, p. 328)

in an ideal society ruled by convention, there is but the divine will or perhaps, by assumption, a common will of all individuals concerning social decisions, so that [...] no conflict of individual wills is involved.

This ‘common will’ is precisely that discussed above in the context of neural and insect populations, with the only complication being that individual voters can hold mistaken views about what is best. Then, ‘the methods of

[...] convention are, or can be, rational in the sense that any individual can be rational in his choice’ (*ibid.*). Thus it is unsurprising that an optimal strategy for aggregating the preferences of individuals, while non-trivial due to the potential for individuals to be mistaken, is still possible. It is interesting to contrast Arrow’s view with that of another foundational thinker in voting theory, Nicolas de Condorcet. Among Condorcet’s ideas (Condorcet 1785), his ‘jury theorem’ has recently been proposed as relevant to the study of collective decision-making in animals (List 2004). This result shows how, for voting by groups with convergent interests in choosing between two alternatives, but where each individual may be mistaken in their assessment of the better option, the probability of the group collectively selecting the better option varies according to its size and the reliability of its members’ information. This is interesting, in that it shows that majority decisions are better than individual decisions when individuals are sufficiently likely to correctly identify the better alternative. List (2004) proposed that Condorcet’s ‘jury theorem’ could be applicable to nest-site selection by honeybees, but to the best of my knowledge this has not been attempted. This could be because the result neglects one crucial aspect of collective decision-making by social insects, interaction between the ‘voters’. While the Condorcet approach does make useful predictions for collective behaviour, considering interactions between voting populations allows stronger claims based on statistical optimality to be made (Bogacz *et al.* 2006; Marshall *et al.* 2009).

Returning to Arrow’s work, despite its apparently being formulated to address a different problem to that facing groups with convergent interests, what does seem interesting is the concept of individual and group rationality. Arrow considers individual rationality in terms of the classical transitivity of preferences; ‘if x is preferred to y and y to z , then x is preferred to z ’, for example (Arrow 1950, p.331). Again, there is no problem with this concept in the systems discussed above; for example, choice transitivity is also observed in ants (Franks *et al.* 2003). Then, according to Arrow, a group-level preference is rational if it faithfully captures the majority preferences of its constituent members (Arrow 1950, p. 329). Here, an intriguing contrast with recent experimental evidence on collective choice in ants presents itself. Another important economic definition of rationality is that a decision-maker be immune to the ‘distractor effect’ in multi-objective decisions (or ‘irrelevance of independent alternatives’, reviewed in Bateson 2010). For decisions between alternatives each having two distinct objective values, for example, the pattern of choice between two alternatives differing in these two objectives should not be affected by the introduction of a third alternative that is inferior to the other two in *both* objectives. This condition is closely

related to Arrow's 'independence of irrelevant alternatives' condition (Arrow 1950, p. 337), although since Arrow's result considers only preference orderings rather than distributions, it cannot be exactly captured in his framework (the relationship between various definitions of 'independence of irrelevant alternatives' is discussed in McLean 1995). Ant colonies have been shown to be rational, in that they are immune to distractor effects (Edwards and Pratt 2009). However, more recently it has been shown that this is *despite* the preferences of individual colony members being *vulnerable* to distractor effects (Sasaki and Pratt 2011). Although individuals are consistent in their vulnerability to distractors, the collective decision-making mechanism that aggregates their preferences into a collective preference 'irons out' these individual irrationalities into a collective decision that is simultaneously rational, in the sense that it is independent of irrelevant alternatives, and irrational (in Arrow's sense), in that the collective choice does not reflect the majority preference ordering of the individual voters. The collective-level immunity to distractor effects can only increase colony fitness. Understanding the mechanism by which this is achieved is an interesting open problem at the current time, but it likely to involve the facts that individuals recruit each other, and that many individuals are observed not to evaluate multiple alternatives (S.C. Pratt, pers. comm.). This also highlights another important difference between voting in the sense of Arrow and Condorcet, and in the sense of populations of neurons, social insects and humans. In both Arrow's and Condorcet's frameworks individual preferences are fixed, and individuals do not seek to influence each other (although Condorcet did consider the need to account for this; see Urken 2008). In more sophisticated systems, individual preferences are fluid, and individuals do seek to influence each other. This fluidity and mutual influence helps explain how optimal voting might in fact be possible, in groups where the members' interests converge.

Acknowledgements

I thank A.B. Urken and S.C. Pratt for helpful comments and discussions.

References

Arrow, K.J. 1950. A difficulty in the concept of social welfare. *The Journal of Political Economy* 58: 328-346.
 Bateson, M. 2010. Rational choice behaviour: definitions and evidence. In *Encyclopedia of Animal Behaviour* 3: 13-19. Oxford: Academic Press.
 Bogacz, R., Brown, E., Moehlis, J., Holmes, P. and Cohen, J.D. 2006. The physics of optimal decision making: a

formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review* 113: 700-765.
 Condorcet, J.A.N. 1785. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Paris.
 Edwards, S.C. and Pratt, S.C. 2009. Rationality in collective decision-making by ant colonies. *Proceedings of the Royal Society of London B*, 276: 3655-3661.
 Franks, N.R., Mallon, E.B., Bray, H.E., Hamilton, M.J. and Mischler, T.C. 2003. Strategies for choosing between alternatives with different attributes: exemplified by house-hunting ants. *Animal Behavior* 65: 215-223.
 George, A.J.T., Stark, J. and Cliburn, C. 2005. Understanding specificity and sensitivity of T-cell recognition. *Trends in Immunology* 26: 653-659.
 Gittins, J.C. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society. Series B (Methodological)* 41: 148-177.
 List, C. 2004. Democracy in animal groups: a political science perspective. *Trends in Ecology and Evolution* 19: 168-169.
 Marshall, J.A.R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T. and Franks, N.R. 2009. On optimal decision making in brains and social insect colonies. *Journal of the Royal Society: Interface* 6: 1065-1074.
 Marshall, J.A.R. and Franks, N.R. 2009. Colony-level cognition. *Current Biology* 19: R395-R396.
 McLean, I. 1995. Independence of irrelevant alternatives before Arrow. *Mathematical Social Sciences* 30: 107-126.
 Möglich, M. 1978. Social organization of nest emigration in *Leptothorax* (Hym., Form.). *Insectes Sociaux* 25, 205-225.
 Perkins, T.J. and Swain, P.S. 2009. Strategies for cellular decision-making. *Molecular Systems Biology* 5: 1-15.
 Ratcliff, R. 1978. A theory of memory retrieval. *Psychological Review* 85: 59-108.
 Sasaki, T. and Pratt, S.C. 2011. Emergence of group rationality from irrational individuals. *Behavioral Ecology and Sociobiology* (in press) doi:10.1093/beheco/arq198.
 Shadlen, M.N. and Newsome, W.T. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology* 86: 1916-1936.
 Seeley, T.D. 2010. *Honeybee Democracy*. Princeton, NJ.: Princeton University Press.
 Urken, A.B. 2008. Grokking Condorcet's 1785 Essai. *Electronic Journal for History of Probability and Statistics* 4: 1-18.
 Visscher, P.K. 2007. Group decision making in nest-site selection among social insects. *Annual Reviews of Entomology* 52: 255-275.
 von Frisch, K. 1967. *The Dance Language and Orientation of Bees*. Cambridge, Mass.: Belknap Press.