

Computational Persuasion with Applications in Behaviour Change

Anthony HUNTER

*Department of Computer Science,
University College London,
Gower Street, London WC1E 6BT, UK
anthony.hunter@ucl.ac.uk*

Abstract. Persuasion is an activity that involves one party trying to induce another party to believe something or to do something. It is an important and multifaceted human facility. Obviously, sales and marketing is heavily dependent on persuasion. But many other activities involve persuasion such as a doctor persuading a patient to drink less alcohol, a road safety expert persuading drivers to not text while driving, or an online safety expert persuading users of social media sites to not reveal too much personal information online. As computing becomes involved in every sphere of life, so too is persuasion a target for applying computer-based solutions. An automated persuasion system (APS) is a system that can engage in a dialogue with a user (the persuadee) in order to persuade the persuadee to do (or not do) some action or to believe (or not believe) something. To do this, an APS aims to use convincing arguments in order to persuade the persuadee. Computational persuasion is the study of formal models of dialogues involving arguments and counterarguments, of user models, and strategies, for APSs. A promising application area for computational persuasion is in behaviour change. Within healthcare organizations, government agencies, and non-governmental agencies, there is much interest in changing behaviour of particular groups of people away from actions that are harmful to themselves and/or to others around them.

Keywords. Computational persuasion; Persuasion dialogues; Persuasive arguments; Dialogical argumentation; Computational models of argument; Probabilistic argumentation; Argumentation strategies.

1. Introduction

Persuasion is an activity that involves one party trying to get another party to do (or not do) some action or to believe (or not believe) something. It is an important and multifaceted human facility. Consider, for example, a doctor persuading a patient to drink less, a road safety expert persuading drivers to not text while driving, or an online safety expert persuading users of social media sites to not reveal too much personal information.

In this paper, I discuss some aspects of the notion of persuasion, and explain how this leads to the idea of computational persuasion. Computational models of

argument are central to the development of computational persuasion. I briefly review some key aspects of computational models of argument, and highlight some topics that need further development. I then briefly cover behaviour change as a topic that we can apply methods from computational persuasion, and evaluate the progress in the field.

2. What is persuasion?

The aim of persuasion is for the persuader to change the mind of the persuadee. Some kinds of interaction surrounding persuasion include: Persuader collecting information, preferences, etc from the persuadee; Persuader providing information, offers, etc to the persuadee; Persuader winning favour (e.g. by flattering the persuadee, by making small talk, by being humorous, etc); But importantly, arguments are the essential structures for presenting the claims (and counter claims) in persuasion. An argument-centric focus on persuasion leads to a number of inter-related aspects (see list below) that need to be taken into account, any of which can be important in bringing about successful persuasion.

Persuader The nature of the persuader can be important. From a rational perspective, seemingly good features of a persuader are that s/he has relevant authority, expertise, or knowledge, and seemingly poor features of a persuader are that s/he is attractive, witty, or a celebrity. However, in practice, different persuadees respond to different features. For instance, a teenager is unlikely to be convinced by a government safety expert to wear a helmet when on a bike, but may be influenced by a celebrity to do so.

Language The choice of language in argumentation can be important. This goes from from choice of words (e.g. use of *freedom fighter* versus *terrorist*), to choice of metaphor, or use of irony[17].

Psychology The use of psychological techniques can be important [16] such as: Reciprocation (e.g. doing a small favour for someone is more likely to result in a big favour being obtained in return); Consistency (e.g. getting expressed support for a cause, prior to asking for material support is more likely to be successful); And social proof (e.g. treating dog phobia in children by showing videos of children playing happily with children).

Personality Determining the personality of the persuadee can be important. Consider for example persuading someone to vote in the national election: If the person “follows the crowd”, then tell them that the majority of the population voted in the last election, whereas if the person “follows rules rigorously”, then tell them that it is their duty to vote. Mistaking the personality trait can have a negative effect on the chances of successful persuasion.

Rationality Presenting rational arguments can be important. If a persuader wants to convince the persuadee of an argument (a persuasion argument), then this includes acceptability of the persuasion argument (against counterarguments), believing the premises of the persuasion argument, fit of persuasion argument with agenda, goals, preferences, etc, quality of constellation of arguments considered (balance, depth, breadth, understandability, etc).

Emotion Presenting emotional arguments can be important. For example, *you have a good income, and so you should feel guilty if you do not donate money to this emergency appeal by Médecins Sans Frontières*. As another example, *your parents will be proud of you if you complete your thesis and get your PhD award*. Note, emotional arguments contrast with evidential/logical arguments (e.g. *You will have a much higher chance of getting a highly paid job if you complete your thesis and get your PhD award*).

The above dimensions that can affect the success of argumentation can be considered together in the following criterion for successful persuasion.

Selectivity Persuasion does not involve exhaustive presentation of all possible arguments [8]. Rather it requires careful selection of arguments that are most likely to be efficacious in changing the mind of the persuadee. Deciding on which arguments to select depends on diverse features of the arguments and the persuadee such as the nature of the persuader, the language of the arguments, use of psychological techniques, personality of the persuadee, use of rational and/or emotional argumentation, etc.

Being selective does not mean that argumentation needs to be constrained in any way other than being the most efficacious for persuasion. In particular, I would like to make the following claim.

Persuasion is not normative There are no underlying rules or principles to the use of argumentation in persuasion. This means for instance that arguments can be inconsistent, irrational, untrue, etc. if they persuade. Though inconsistent, irrational, untrue arguments may be counter-productive with some audiences, as well as being potentially problematic from moral, ethical, and regulatory perspectives.

A corollary of the above claim is that how convincing an argument is does not equal how correct it is. For example, arguments like *homeopathy focuses on processes of health and illness rather than states, and therefore it is better than regular medicine* and *the sheer weight of anecdotal evidence gives rise to the common-sense notion that there must be some basis for homeopathic therapies by virtue of the fact that they have lasted this long* can be convincing for some audiences.

3. What is computational persuasion?

An **automated persuasion system** (APS), i.e. a persuader, is a system that can engage in a dialogue with a user, i.e. a persuadee, in order to persuade that persuadee to do (or not do) some action or to believe (or not believe) something. To do this, an APS aims to use convincing arguments in order to persuade the persuadee. The dialogue may involve moves including queries, claims, and importantly, arguments and counterarguments, that are presented according to some protocol. Whether an argument is convincing depends on the context, and on the characteristics of the persuadee. An APS maintains a model of the persuadee,

and this is harnessed by the strategy of the APS in order to choose good moves to make in the dialogue.

Computational persuasion is the study of formal models of dialogues involving arguments and counterarguments, of persuadee models, and strategies, for APSs. Therefore, developments in computational persuasion build on computational models of argument. Note, the aim of computational persuasion is not to produce models of human persuasion (c.f. [11]), rather it is to produce models of persuasion that can be used by computers to persuade humans, and that they can be shown to have a reasonable success rate in some persuasion goal (i.e. that a reasonable proportion of the users are persuaded by the arguments and therefore do the action or accept the belief).

3.1. What do computational models of argument offer?

Computational persuasion is based on computational models of argument. These models are being developed to reflect aspects of how humans use conflicting information by constructing and analyzing arguments. A number of models have been developed, and some basic principles established. We can group much of this work in four levels as follows (with only examples of relevant citations).

Dialectical level Dialectics is concerned with determining which arguments win in some sense. In abstract argumentation, originally proposed in the seminal work by Dung [23], arguments and counterarguments can be represented by a graph. Each node denotes an argument, and each arc denotes one argument attacking another argument. Dung defined some principled ways to identify extensions of an argument graph. Each extension is a subset of arguments that together act as a coalition against attacks by other arguments. An argument in an extension is, in a sense, acceptable. Methods for argument dynamics ensure that specific arguments hold in the extensions of the argument graph such as epistemic enforcement in abstract argumentation [4,3,18], revision of argument graphs [19,20], and belief revision in argumentation (e.g. [14,27,10,22]).

Logical level At the dialectic level, arguments are atomic. They are assumed to exist, but there is no mechanism for constructing them. Furthermore, they cannot be divided or combined. To address this, the logical level provides a way to construct arguments from knowledge. At the logical level, an argument is normally defined as a pair $\langle \Phi, \alpha \rangle$ where Φ is a minimal consistent subset of the knowledgebase (a set of formulae) that entails α (a formula). Here, Φ is called the support, and α is the claim, of the argument. Hence, starting with a set of formulae, arguments and counterarguments can be generated, where a counterargument (an argument that attacks another argument) either rebuts (i.e. negates the claim of the argument) or undercuts (i.e. negates the support of the argument). A range of options for structured argumentation at the logic level have been investigated (see [9,61,64,28] for tutorial reviews of some of the key proposals).

Dialogue level Dialogical argumentation involves agents exchanging arguments in activities such as discussion, debate, persuasion, and negotiation. Starting with [31,43], dialogue games are now a common approach to characterizing argumentation-based agent dialogues (e.g. [1,12,21,24,45,46,50,51,65]). Dialogue games are normally made up of a set of communicative acts called moves, and a

protocol specifying which moves can be made at each step of the dialogue. Dialogical argumentation can be viewed as incorporating logic-based argumentation, but in addition, dialogical argumentation involves representing and managing the locutions exchanged between the agents involved in the argumentation. The emphasis of the dialogical view is on the interactions between the agents, and on the process of building up, and analyzing, the set of arguments until the agents reach a conclusion. See [52] for a review of formal models of persuasion dialogues and [62,13] for reviews and analyses of strategies in dialogical argumentation.

Rhetorical level Normally argumentation is undertaken in some wider context of goals for the agents involved, and so individual arguments are presented with some wider aim. For instance, if an agent is trying to persuade another agent to do something, then it is likely that some rhetorical device is harnessed and this will affect the nature of the arguments used (e.g. a politician may refer to *investing in the future of the nation's children* as a way of persuading colleagues to vote for an increase in taxation). Aspects of the rhetorical level include believability of arguments from the perspective of the audience [32], impact of arguments from the perspective of the audience [33], use of threats and rewards [2], appropriateness of advocates [34], and values of the audience [5,6,48].

So computational models of argument offer a range of formal systems for generating and comparing arguments, and for undertaking this in a dialogue.

3.2. Shortcomings in the state of the art

However there are shortcomings in the state of the art of computational models of argument for application in persuasion. The current state of the literature does not adequately offer the following and hence there are some exciting research challenges to be addressed if we are to deliver computational persuasion.

Domain knowledge A formalization of domain knowledge appropriate for constructing arguments concerning behaviour change (e.g. a formalism for representing persuadee preferences, persuadee goals, persuadee preferences, system persuasion goals, and system knowledge concerning actions that can address persuadee goals, etc) though the multiagent communities offer proposals that might be adapted for our needs.

Persuasion protocols Protocols that take account of humans unable to make rich input (since we are not supporting free text input from the persuadee).

Persuadee models Persuadee models that allow the persuasion system to construct a model of the persuadee's beliefs and preferences, to qualify the probabilistic uncertainty of that model, and to update that model and the associated uncertainty as the dialogue progresses, though some promising proposals could contribute to our solution (e.g. [29,36,58,38]).

Persuasion strategies Strategies for persuasion that harness the persuadee model to find optimal moves to make at each stage (trading the increase in probability of successfully persuading the persuadee against the raised risk that the persuadee disengages from the dialogue as it progresses).

In order to focus research on addressing these shortcomings, we can consider how computational persuasion can be developed and evaluated in the context of behaviour change applications.

| Field | Examples of behaviour change topic |
|-----------------------|--|
| Healthy life-styles | eating fewer calories, eating more fruit and veg, doing more exercise, drinking less alcohol |
| Addiction management | gambling, smoking, drugs |
| Treatment compliance | self-management of diabetes, taking vaccines, completing course of antibiotics |
| Personal finance | borrowing less, saving more |
| Education | starting or continuing with a course, studying properly |
| Energy efficiency | reducing electricity consumption, installing home insulation |
| Citizenship | voting, recycling, contributing to charities, wasting less food |
| Safe driving | not exceeding speed limits, not texting while driving |
| Anti-social behaviour | aggression, vandalism, racism, sexism, trolling |

Table 1. Some examples where people could change their behaviour and for which there would be a substantial quantifiable benefit to themselves, and/or to society.

4. What is behaviour change?

There is a wide variety of problems that are dangerous or unhealthy or unhelpful for an individual, or for those around them, and that are expensive to government and/or to society (see Table 1 for examples). For each type of problem, we can conceivably tackle a small proportion of cases with substantial benefit to individuals, government and society using techniques for behaviour change.

Many organizations are involved in behaviour change, and many approaches are used to persuade people to change their behaviour including counselling, information resources, and advertising. Many diverse factors can influence how such approaches can be used effectively in practice such as the following.

- Perceived social norms (e.g. everyone drives above the speed limit).
- Social pressure (e.g. my friends laugh at me if I drive slowly).
- Emotional issues (e.g. speeding is cool).
- Agenda (e.g. I am always late for everything, and so I have to speed).
- Perception of an issue (e.g. I am a good driver even if I speed).
- Opportunities to change behaviour (e.g. access to a race track on which to drive fast instead of driving fast on ordinary roads).
- Attitude to persuader (e.g. I listen to Lewis Hamilton not a civil servant).
- Attitude to information (e.g. I switch off if I am given statistics).

As computing becomes involved in every sphere of life, so too is persuasion a target for applying computer-based solutions. There are **persuasion technologies** that have come out of developments in human-computer interaction research (see for example the influential work by Fogg [26]) with a particular emphasis on addressing the need for systems to help people make positive changes to their behaviour, particularly in healthcare and healthy life-styles.

Many of these persuasion technologies for behaviour change are based on some combination of questionnaires for finding out information from users, provision of information for directing the users to better behaviour, computer games to enable users to explore different scenarios concerning their behaviour, provision

of diaries for getting users to record ongoing behaviour, and messages to remind the persuadee to continue with the better behaviour.

Interestingly, argumentation is not central to the current manifestations of persuasion technologies. The arguments for good behaviour seem either to be assumed before the persuadee accesses the persuasion technology (e.g. when using diaries, or receiving email reminders), or arguments are provided implicitly in the persuasion technology (e.g. through provision of information, or through game playing). So explicit consideration of arguments and counterarguments are not supported with existing persuasion technologies. This creates interesting opportunities for computational persuasion to develop APSs for behaviour change where arguments are central.

5. How can computational persuasion be applied?

Computational models of argument drawing on ideas of abstract argumentation, logical argumentation, dialogical argumentation, together with techniques for argument dynamics and for rhetorics, offer an excellent starting point for developing computational persuasion for applications in behaviour change.

I assume that an APS for behaviour change is a software application running on a desktop or mobile device. Some difficult challenges to automate persuasion via an app are the following.

1. Need asymmetric dialogues without natural language interface.
2. Need short dialogues to keep engagement.
3. Need well-chosen arguments to maximize impact.
4. Need to model the user in order to be able to optimize the dialogue.
5. Need to learn from previous interactions with the agent or similar agents.
6. Need to model the domain to generate arguments/counterarguments.

The dialogue may involve steps where the system finds out more about the persuadee's beliefs, intentions and desires, and where the system offers arguments with the aim of changing the persuadee's beliefs, intentions and desires. The system also needs to handle objections or doubts (represented by counterarguments) with the aim of providing a dialectically winning position. To illustrate how a dialogue can lead to the presentation of an appropriate context-sensitive argument consider the example in Table 2. In this, only the APS presents arguments, and when it is the user's turn s/he can only answer questions (e.g. yes/no questions) or select arguments from a menu. In Figure 1, a dialogue step is illustrated where a user can state the degree of agreement or disagreement in an argument.

Arguments can be automatically generated from a knowledgebase. For this, we can build a knowledgebase for each domain, though there are many commonalities in the knowledge required for each behaviour change application.

- Persuadee beliefs (e.g. cakes give a sugar rush).
- Persuadee preferences (e.g. burgers are preferred to apples).
- Behavioural states (e.g. persuadee's weight, exercise regime, etc.).
- Behavioural actions (e.g. eat a piece of fruit, eat a piece of cake, walk 1km).
- Behavioural goals (e.g. lose 10Kg by Christmas, reduce sugar intake).

| Step | Who | Move |
|------|------|---|
| 1 | APS | To improve your health, you could join an exercise class |
| 2 | User | Exercise classes are boring |
| 3 | APS | For exciting exercise, you could do an indoor climbing course |
| 4 | User | It is too expensive |
| 5 | APS | Do you work? |
| 6 | User | No |
| 7 | APS | If you are registered unemployed, then the local sports centre offers a free indoor climbing course |
| 8 | APS | Would you try this? |
| 9 | User | Yes |

Table 2. Simple example of an asymmetric dialogue between a user and an APS. As no natural language processing is assumed, the arguments posted by the user are actually selected by the user from a menu provided by the APS.

Figure 1. Interface for an asymmetric dialogue move for asking the user's belief in an argument. The top argument is by the APS, and the second argument is a counterargument presented by the APS. The user uses the menu to give his/her belief in the counterargument.

To represent and reason with the domain knowledge, we can harness a form of BDI calculus in predicate logic for relating beliefs, behavioural goals, and behavioural states, to possible actions. We can then use the calculus with logical argumentation to generate arguments for persuasion. A small example of an argument graph that we might want to generate by this process is given in Figure 2 including the persuasion goal *giving up smoking will be good for your health*.

To support the selection of arguments, we require persuadee models. For this,

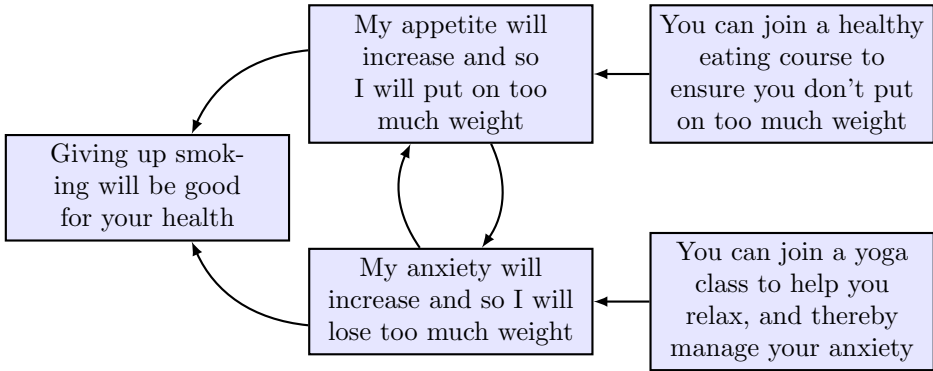


Figure 2. Example of an argument graph for persuasion

we can establish the probabilistic uncertainty associated with the APS model of the persuadee’s beliefs, behavioural state, behavioural goals, preferences, and tendencies etc by asking the persuadee appropriate questions, by considering previous usage of the APS by the persuadee, and by the general class of the persuadee (i.e. by assignment to a built-in model learned from a class of similar users). Key possible dimensions for modelling uncertainty are given in Table 3.

Two main approaches to probabilistic argumentation are the constellations and the epistemic approaches [37].

- In the constellations approach, the uncertainty is in the topology of the graph (see for example [42,35]). As an example, this approach is useful when one agent is not sure what arguments and attacks another agent is aware of, and so this can be captured by a probability distribution over the space of possible argument graphs.
- In the epistemic approach, the topology of the argument graph is fixed, but there is uncertainty about whether an argument is believed [63,37,41]. A core idea of the epistemic approach is that the more likely it is to believe in an argument, the less likely it is to believe in an argument attacking it. The epistemic approach can give a finer grained version of Dung’s approach, and it can be used to give a valuable alternative to Dung’s approach. For example, for a graph containing arguments A and B where B attacks A , it might be the case that a user believes A and not B , and if so the epistemic extension (the set of believed arguments) would be $\{A\}$ which is in contrast the Dung’s approach where the only extension is $\{B\}$.

There are approaches to bringing probability theory into systems for dialogical argumentation. A probabilistic model of the opponent has been used in a dialogue strategy allowing the selection of moves for an agent based on what it believes the other agent is aware of [57]. In another approach to probabilistic opponent modelling, the history of previous dialogues is used to predict the arguments that an opponent might put forward [29]. Though further avenues need to be explored.

The constellations approach can model the uncertainty about the structure of the graph in the persuadee mind. We can update the model with each argu-

| Type of uncertainty | Modelling technique |
|--------------------------------------|-------------------------|
| Beliefs of persuadee | Epistemic approach |
| Arguments/attacks known by persuadee | Constellations approach |
| Moves that persuadee makes | PFSMs/POMDPs |
| Risk of disengagement | Markov models |

Table 3. Possible dimensions of uncertainty in models of persuadee

ment/attack presented. Also, we can use expected utility to identify best choice of argument/attack to present [40].

The epistemic approach is useful for asymmetric dialogues where the user is not allowed to posit arguments or counterarguments [39]. So the only way the user can treat arguments that s/he does not accept is by disbelieving them. In contrast, in symmetric dialogues, the user could be allowed to posit counterarguments to an argument that s/he does not accept. The distribution can be updated in response to moves made (posits, answers to queries, etc) using different assumptions about the persuadee (credulous, skeptical, rational, etc). The aim is to choose moves that will increase belief in positive persuasion goals or decrease belief in negative persuasion goals.

For modelling the possible dialogues that might be generated by a pair of agents, a probabilistic finite state machine can represent the possible moves that each agent can make in each state of the dialogue assuming a set of arguments that each agent is aware of [38]. Each state is composed of the public state of the dialogue (e.g. what has been said) and the private state of each participant (e.g. the arguments they believe). We can find optimal sequences of moves by handling uncertainty concerning the persuadee using partially observable markov decision processes (POMDPs) when there is uncertainty about the private state of the persuader [30].

A strategy for an APS needs to find the best choice of move at each stage where best is determined in terms of some combination of the need to increase the likelihood that the persuadee is persuaded by the goal of the persuasion, and the need to decrease the likelihood that the persuadee disengages from the dialogue. For instance, at a certain point in the dialogue, the APS might have a choice of two arguments A and B to present. Suppose A involves further moves to be made (e.g. supporting arguments) whereas B is a single posit. So choosing A requires a longer dialogue (and higher probability of disengagement) than B . However, if the persuadee keeps to the end of each dialogue, then it is more likely that the persuadee believes A than B . An APS should present arguments and counterarguments that are informative, relevant, and believable, to the persuadee. If the APS presents uninformative, irrelevant, or unbelievable arguments (from the perspective of the persuadee), the probability of successful persuasion is reduced, and it may alienate the persuadee. A choice of strategy depends on the protocol, and on the kind of dynamic persuadee model. Various parameters can be considered in the strategy such as the preferences of the persuadee, the agenda of the persuadee, etc.

Probabilistic models of the opponent have been used in some strategies allowing the selection of moves for an agent based on what it believes the other

agent believes [36]. Utility theory has also been considered in argumentation (for example [54,59,44,49]) though none of these represents the uncertainty of moves made by each agent in argumentation. Probability theory and utility theory (using decision theory) has been used in [40] to identify outcomes with maximum expected utility where outcomes are specified as particular arguments being included or excluded from extensions. Strategies in argumentation have also been analyzed using game theory [53,55,25], though these are more concerned with issues of manipulation, rather than persuasion.

Given that we need to consider multiple dimensions in identifying a more convincing argument (e.g. whether an argument is believed, whether an argument is undefeated, whether it is relevant, whether it relates to the goals of the persuadee, etc), there is a need to generalize the existing proposals for strategies for argumentation.

6. Discussion

Computational persuasion, being based on computational models of argument, is a promising approach to technology for behaviour change applications. Developing an APS involves research challenges including: undertaking the dialogue without using natural language processing; having an appropriate model of the domain in order to identify arguments; having an appropriate dynamic model of the persuadee; and having a strategy that increases the probability of persuading the persuadee. Furthermore, with even a modest set of arguments, the set of possible dialogues can be enormous, and so the protocols, persuadee models, and strategies need to be computationally viable.

In the short-term, we may envisage that the dialogues between an APS and a user involve limited kinds of interaction. For example, the APS manages the dialogue by asking queries of the persuadee, where the allowed answers are given by a menu or are of restricted types (e.g. age), and by posing arguments, and the persuadee may present arguments that are selected from a menu presented by the APS. Obviously richer natural language interaction would be desirable, but it is not feasible in the short-term. Even with such restricted asymmetric dialogues, it may be possible that effective persuasion can be undertaken, and furthermore, we need to investigate this conjecture empirically with participants.

There are some investigations of computational models of argument with participants. In a study by Rahwan *et al* [56], participants were given argument graphs and asked about their confidence in specific arguments being acceptable or unacceptable. Interestingly, for an unattacked argument A that is then attacked by a new argument B , the confidence in A being acceptable does not necessarily fall to zero (as would be predicted by the usual dialectical semantics for abstract argumentation). Then if a further new argument C is added that attacks B , the confidence in A being acceptable does not necessarily rise to 1 (as would be predicted by the usual dialectical semantics for abstract argumentation). In another study, Cerutti *et al* [15], investigated how well an approach to structured argumentation by Prakken and Sartor models how a group of participants reason with three different argumentation scenarios. Their results showed that a corre-

spondence between the acceptability of arguments by participants and the justification status predicated by the structured argumentation in the majority of the cases. But in some cases, the implicit knowledge about domains could substantially affect this. In a study of argumentation dialogues, Rosenfeld and Kraus [60] undertook studies with participants in order to develop a machine learning-based approach to predict the next move a participant would make in a dialogue. Emotion in argumentation has also been the subject of a study with participants in a debate where the emotional state was estimated from EEG data and automated facial expression analysis. In this study, Benlamine *et al* [7] showed for instance that the number and the strength of arguments, attacks and supports exchanged between a participant could be correlated with particular emotions of the participant. There are also relevant studies investigating the efficacy of using arguments as a way of persuading people when compared with other counselling methods indicating that argumentation may have disadvantages if used inappropriately [47]. Whilst these studies only consider some aspects of computational models of argument, they point to the need for further studies with participants if we are to develop a well-understood and well-grounded framework for computational persuasion.

Acknowledgements

This research is part-funded by EPSRC grant EP/N008294/1 *Framework for Computational Persuasion*¹

References

- [1] L. Amgoud, N. Maudet, and S. Parsons. Arguments, dialogue and negotiation. In *Proceedings of ECAI'00*, pages 338–342. IOS Press, 2000.
- [2] L. Amgoud and H. Prade. Formal handling of threats and rewards in a negotiation dialogue. In *Proceedings of AAMAS'05*, pages 529–536, 2005.
- [3] R. Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *Proc. of ECAI'12*, pages 127–132, 2012.
- [4] R. Baumann and G. Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proc. of COMMA'10*, pages 75–86, 2010.
- [5] T. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
- [6] T. Bench-Capon, S. Doutre, and P. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.
- [7] S. Benlamine, M. Chaouachi, S. Villata, E. Cabrio, C. Frasson, and F. Gandon. Emotions in argumentation: an empirical evaluation. In *Proc. of IJCAI'15*, pages 156–163, 2015.
- [8] Ph. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.
- [9] Ph. Besnard and A. Hunter. Constructing argument graphs with deductive arguments: a tutorial. *Argument and Computation*, 5(1):5–30, 2014.
- [10] P. Bisquert, C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasquie-Schiex. Enforcement in argumentation is a kind of update. In *Proc. of SUM'13*, volume 8078 of *LNCS*, pages 30–42. Springer, 2013.
- [11] P. Bisquert, M. Croitoru, and F. Dupin de Saint-Cyr. Four ways to evaluate arguments according to agent engagement. In *Proc. of Brain Informatics and Health*, volume 9250 of *LNCS*. Springer, 2015.

¹For more information on the project, see www.computationalpersuasion.com.

- [12] E. Black and A. Hunter. An inquiry dialogue system. *Autonomous Agents and Multi-Agent Systems*, 19(2):173–209, 2009.
- [13] E. Black and A. Hunter. Reasons and options for updating an opponent model in persuasion dialogues. In *Proc. of TAFA'15*, volume 9524 of *LNCS*, pages 21–39. Springer, 2015.
- [14] C. Cayrol, F. Dupin de Saint-Cyr, and M.-C. Lagasque-Schiex. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, 38:49–84, 2010.
- [15] F. Cerutti, N. Tintarev, and N. Oren. Formal arguments, preferences, and natural language interfaces to humans: an empirical evaluation. In *Proc. of ECAI'14*, pages 207–212, 2014.
- [16] R. Cialdini. *Influence: The Psychology of Persuasion*. HarperCollins, 1984.
- [17] R. Cockcroft and S. Cockcroft. *Persuading People*. Macmillan, 1992.
- [18] S. Coste-Marquis, S. Konieczny, and J-G. Maily. Extension enforcement in abstract argumentation as an optimization problem. In *Proc. of IJCAI'15*, pages 2876–2882, 2014.
- [19] S. Coste-Marquis, S. Konieczny, and J-G. Maily. On the revision of argumentation systems: Minimal change of argument statuses. In *Proc. of KR'14*, pages 72–81, 2014.
- [20] S. Coste-Marquis, S. Konieczny, and J-G. Maily. A translation-based approach for revision of argumentation frameworks. In *Proc. of JELIA'14*, pages 77–85, 2014.
- [21] F. Dignum, B. Dunin-Keplicz, and R. Verbrugge. Dialogue in team formation. In *Issues in Agent Communication*, pages 264–280. Springer, 2000.
- [22] M. Diller, A. Haret, T. Linsbichler, S. Rümmele, and S. Woltran. An extension-based approach to belief revision in abstract argumentation. In *Proc. of IJCAI'15*, pages 2926–2932, 2015.
- [23] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- [24] X. Fan and F. Toni. Assumption-based argumentation dialogues. In *Proc. of IJCAI'11*, pages 198–203, 2011.
- [25] X. Fan and F. Toni. Mechanism design for argumentation-based persuasion. In *Proc. of COMMA'12*, pages 322–333, 2012.
- [26] B. Fogg. Persuasive computers. In *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, pages 225–232. CHI, 1998.
- [27] D. Gabbay and O. Rodrigues. A numerical approach to the merging of argumentation networks. In *Proc. of CLIMA'12*, volume 7486 of *LNCS*, pages 195–212. Springer, 2012.
- [28] A. Garcia and G. Simari. Defeasible logic programming: Delp-servers, contextual queries, and explanations for answers. *Argument and Computation*, 5(1):63–88, 2014.
- [29] C. Hadjinikolis, Y. Siantos, S. Modgil, E. Black, and P. McBurney. Opponent modelling in persuasion dialogues. In *Proc. of IJCAI'13*, page 164170, 2013.
- [30] E. Hadoux, A. Beynier, N. Maudet, P. Weng, and A. Hunter. Optimization of probabilistic argumentation with markov decision models. In *Proc. of IJCAI'15*, 2015.
- [31] C. Hamblin. Mathematical models of dialogue. *Theoria*, 37:567–583, 1971.
- [32] A. Hunter. Making argumentation more believable. In *Proc. of AAAI'04*, pages 269–274, 2004.
- [33] A. Hunter. Towards higher impact argumentation. In *Proc. of AAAI'04*, pages 275–280, 2004.
- [34] A. Hunter. Reasoning about the appropriateness of proponents for arguments. In *Proc. of AAAI'08*, pages 89–94, 2008.
- [35] A. Hunter. Some foundations for probabilistic abstract argumentation. In *Proc. of COMMA'12*, pages 117–128. IOS Press, 2012.
- [36] A. Hunter. Modelling uncertainty in persuasion,. In *Proc. of SUM'13*, volume 8078 of *LNCS*, pages 57–70. Springer, 2013.
- [37] A. Hunter. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning*, 54(1):47–81, 2013.
- [38] A. Hunter. Probabilistic strategies in dialogical argumentation. In *Proceedings of SUM'14*, volume 8720 of *LNCS*, pages 190–202. Springer, 2014.
- [39] A. Hunter. Modelling the persuadee in asymmetric argumentation dialogues for persua-

- sion. In *Proc. of IJCAI'15*, 2015.
- [40] A. Hunter and M. Thimm. Probabilistic argument graphs for argumentation lotteries. In *Computational Models of Argument (COMMA'14)*, 2014.
- [41] A. Hunter and M. Thimm. Probabilistic argumentation with incomplete information. In *Proc. of ECAI'14*, pages 1033–1034, August 2014.
- [42] H. Li, N. Oren, and T. J. Norman. Probabilistic argumentation frameworks. In *Proc. of TAFE'11*, 2011.
- [43] J. Mackenzie. Question begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.
- [44] P. Matt and F. Toni. A game-theoretic measure of argument strength for abstract argumentation. In *Proceedings of JELIA'08*, volume 5293 of *LNCS*, pages 285–297, 2008.
- [45] P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11:315–334, 2002.
- [46] P. McBurney, R. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems*, 7:235–273, 2003.
- [47] H. Nguyen and J. Masthoff. Designing persuasive dialogue systems: Using argumentation with care. In *Proceedings of Persuasive technology'08*, pages 201–212, 2008.
- [48] N. Oren, K. Atkinson, and H. Li. Group persuasion through uncertain audience modelling. In *Proc. of COMMA'12*, pages 350–357, 2012.
- [49] N. Oren and T. Norman. Arguing using opponent models. In *Proc. of ArgMAS'09*, volume 6057 of *LNCS*, pages 160–174, 2009.
- [50] S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of some formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.
- [51] H. Prakken. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation*, 15(6):1009–1040, 2005.
- [52] H. Prakken. Formal systems for persuasion dialogue. *Knowledge Engineering Review*, 21(2):163–188, 2006.
- [53] I. Rahwan and K. Larson. Mechanism design for abstract argumentation. In *Proc. of AAMAS'08*, pages 1031–1038, 2008.
- [54] I. Rahwan and K. Larson. Pareto optimality in abstract argumentation. In *Proc. of AAAI'08*, 2008.
- [55] I. Rahwan, K. Larson, and F. Tohmé. A characterisation of strategy-proofness for grounded argumentation semantics. In *Proc. of IJCAI'09*, pages 251–256, 2009.
- [56] I. Rahwan, M. Madakkatel, J. Bonnefon, R. Awan, and S. Abdallah. Behavioural experiments for assessing the abstract argumentation semantics of reinstatement. *Cognitive Science*, 34(8):14831502, 2010.
- [57] T. Rienstra. Towards a probabilistic dung-style argumentation system. In *Proc. of Agreement Technologies (AT'12)*, 2012.
- [58] T. Rienstra, M. Thimm, and N. Oren. Opponent models with uncertainty for strategic argumentation. In *Proc. of IJCAI'13*. IJCAI/AAAI, 2013.
- [59] R. Riveret, H. Prakken, A. Rotolo, and G. Sartor. Heuristics in argumentation: A game theory investigation. In *Proc. of COMMA'08*, pages 324–335, 2008.
- [60] A. Rosenfeld and S. Kraus. Providing arguments in discussions based on the prediction of human argumentative behavior. In *Proc. of AAAI'15*, pages 1320–1327, 2015.
- [61] S. Modgil and H. Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument and Computation*, 5(1):31–62, 2014.
- [62] M. Thimm. Strategic argumentation in multi-agent systems. *Kunstliche Intelligenz*, 28:159–168, 2014.
- [63] M. Thimm. A probabilistic semantics for abstract argumentation. In *Proc. of ECAI'12*, August 2012.
- [64] F. Toni. A tutorial on assumption-based argumentation. *Argument and Computation and Computation*, 5(1):89–117, 2014.
- [65] D. Walton and E. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, 1995.