

# Towards Computational Persuasion via Natural Language Argumentation Dialogues<sup>\*</sup>

Anthony Hunter, Lisa Chalaguine, Tomasz Czernuszenko, Emmanuel Hadoux,  
and Sylwia Polberg

Department of Computer Science,  
University College London,  
London, WC1E 6BT, UK

**Abstract.** Computational persuasion aims to capture the human ability to persuade through argumentation for applications such as behaviour change in healthcare (e.g. persuading people to take more exercise or eat more healthily). In this paper, we review research in computational persuasion that incorporates domain modelling (capturing arguments and counterarguments that can appear in a persuasion dialogues), user modelling (capturing the beliefs and concerns of the persuadee), and dialogue strategies (choosing the best moves for the persuader to maximize the chances that the persuadee is persuaded). We discuss evaluation of prototype systems that get the user’s counterarguments by allowing them to select them from a menu. Then we consider how this work might be enhanced by incorporating a natural language interface in the form of an argumentative chatbot.

**Keywords:** Persuasion, Computational models of argument, Chatbots.

## 1 Introduction

Persuasion is an activity that involves one party trying to induce another party to believe or disbelieve something or to do (or not do) something. It is an important and multifaceted human facility. Obviously, it is essential in commerce and politics, but it is equally important in many aspects of daily life. Consider for example, a child asking a parent for a rise in pocket money, a doctor trying to get a patient to enter a smoking cessation programme, a charity volunteer trying to raise funds for a poverty stricken area, or a government advisor trying to get people to avoid revealing personal details online that might be exploited by fraudsters.

Arguments are a crucial part of persuasion. They may be explicit, such as in a political debate, or they may be implicit, such as in an advert. In a dialogue involving persuasion, counterarguments also need to be taken into account. Participants may take turns in the dialogue with each of them presenting arguments,

---

<sup>\*</sup> Supported by the EPSRC Computational Persuasion grant EP/N008294/1, an EPSRC DTA studentship, and a UCL Summer Research Internship.

**Table 1.** Some examples of potential applications of computational persuasion that could be used to encourage and guide people to change behaviour in healthcare

Issue	Examples
Healthy life-styles	eating fewer calories, eating more fruit and veg, doing more exercise, drinking less alcohol.
Treatment compliance	undertaking self-management of diabetes, completing a course of antibiotics, completing a course of prophylactics.
Treatment reduction	using alternatives to painkillers for premenstrual syndrome, not requesting antibiotics for viral infections.
Problem avoidance	taking vaccines, taking malaria prophylactics, using safe sex practice.
Screening	participating in breast cancer screening, participating in cervical smear screening, self-screening for prostate cancer, breast cancer, bowel cancer, and melanoma.

some of which may be counterarguments to previously presented arguments. So the aim of the persuader is to change the mind of the persuadee through this exchange of arguments. Since some arguments may be more effective than others in such a dialogue, it is valuable for the persuader to have an understanding of the persuadee and of what might work better with her.

### 1.1 Persuasion in behaviour change

As computing becomes involved in every sphere of life, so too is persuasion a target for applying computer-based solutions. Persuasion technologies have come out of developments in human-computer interaction research (see, for example, the influential work by Fogg [18]) with a particular emphasis on addressing the need for systems to help people make positive changes to their behaviour, particularly in healthcare and lifestyle choices. In recent years, a wide variety of systems has been developed to help users to control body weight, reduce fizzy drink consumption, increase physical exercise, and reduce speeding.

Interestingly, explicit use of argumentation is not central to most current manifestations of persuasion technologies. Either arguments are provided implicitly in the persuasion technology (e.g., through provision of information, or through game playing), or it is assumed that the user has considered the arguments for changing behaviour prior to accessing the persuasion technology (e.g., when using diaries, or receiving email reminders). Explicit argumentation with consideration of arguments and counterarguments is not supported with existing persuasion technologies. Yet, for some tasks in behaviour change, an argument-based approach could be highly beneficial, particularly when someone is lacking some key information, and/or entertaining misconceptions about a topic.

This creates some interesting opportunities for artificial intelligence, using computational models of argument, to develop persuasion technologies for be-

**Table 2.** Simple example of a dialogue between a user and an automated persuasion system (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.

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

behaviour change where arguments are central. For reviews of computational models of argumentation, a.k.a computational argumentation, see [6, 52, 3, 5]. Computational models of argument are beginning to offer ways to formalize aspects of persuasion, and with some adaptation and development, they have the potential to be incorporated into computational persuasion tools for behaviour change. For potential applications, see Table 1.

## 1.2 Automated persuasion systems

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 accept a specific argument (the persuasion goal) that encapsulates the reason for a change of behaviour in some specific respect [41]. For example, the persuasion goal might be that the user needs to eat fruit in order to be more healthy, and the system presents supporting arguments (based on evidence, expert opinion, explanation of the fit with the user’s goals, etc.) and counter-arguments to correct misconceptions or inconsistencies in the user’s opinions. To do this, an APS aims to use convincing arguments in order to persuade the persuadee.

Whether an argument is convincing depends on the context, and on the characteristics of the persuadee. An APS maintains a model of the persuadee to predict what arguments and counterarguments the persuadee knows about and/or believes, and this can be harnessed by the strategy of the APS in order to choose good moves to make in the dialogue.

There have already been some promising studies that indicate the potential of using automated dialogues in behaviour change such as using dialogue games for

health promotion [19, 12, 21, 20], conversational agents for encouraging exercise [44, 7] and for promoting plant-based diets [60], dialogue management for persuasion [1], persuasion techniques for healthy eating messages [58], and tailored assistive living systems for encouraging exercise [22]. However, none of these studies have provided a framework that integrates domain modelling and user modelling for strategic argumentation in behaviour change. In the next section, we review a specific framework that addresses these issues.

## 2 Framework for computational persuasion

In order to provide a framework for computational persuasion, we assume an APS has a domain model, a user model, and a dialogue engine, as components and that these are used by the system to enter into a persuasion dialogue with the user. We will explain these components in more detail below.

In addition, in this section, we assume that the interface for an APS does not accept natural language input from the user. Rather, the system provides a menu of counterarguments, and the user selects those that s/he subscribes to. This therefore avoids the problems of natural language processing. We consider how we may drop this restriction in Section 4 by harnessing a simple natural language interface.

### 2.1 Domain modelling

The domain model contains the arguments that can be presented in the dialogue by the system, and it also contains the arguments that the user may entertain. The domain model can be represented by a bipolar argument graph [13]. This is a graph where each node is an argument, and each arc denotes a relationship between pairs of arguments. We consider two types of relationship for an arc from A to B. The first is an attack relationship, and so the arc from A to B denotes that A attacks B (i.e., A is a counterargument for B). The second is a support relationship, and so the arc from A to B denotes that A supports B (i.e., A provides further information that supports for B).

In order to have good quality dialogues, it is important that the argument graph has sufficient depth and breadth of coverage of the topic. Each argument is represented by a premise and claim in a natural language statement. The choice of language may be important for particular audiences. The argument graph also needs to have sufficient depth so that the dialogue can proceed with more than one or two exchanges of argument per participant. The arguments for the argument graph can be obtained from literature of the domain. For example, for healthy eating, there is a large medical literature on arguments about healthy eating. However, arguments that the user may wish to play are often more difficult to obtain. For instance, it is more difficult to find argument for not having a healthy diet. Hence, we have investigated various techniques for acquiring argument using crowdsourcing [15, 16] and for identifying arguments for behaviour change applications based on for example barriers to change that individuals may perceive [14].

## 2.2 User modelling

The user model contains information about the user that can be used by the system for making good choices of move. The information in the user model is what the system believes is true about the user. The key dimensions that we have considered are belief and concerns associated with arguments by users.

**Beliefs** Arguments are formed from premises and a claim, either of which may be explicit or partially implicit. An agent can express a belief in an argument based on the agent’s belief in the premises being true, the claim being implied by the premises, and the claim being true. There is substantial evidence in the behaviour change literature that shows the importance of the beliefs of a persuadee in affecting the likelihood that a persuasion attempt is successful (see for example the review by Ogden [45]). Furthermore, beliefs can be used as a proxy for fine-grained argument acceptability, the need for which was highlighted by empirical studies conducted in [51, 47].

**Concerns** Arguments are statements that contain information about the agent and/or the world. Furthermore, they can refer to impacts on the agent and/or the world. These impacts may relate to the concerns of the agent. In other words, some arguments may have significant impacts on what the agent is concerned about. We associate concerns with arguments, and then for a user model, we obtain or predict the user’s preferences over the concerns.

To illustrate how beliefs (respectively concerns) arise in argumentation, and how they can be harnessed to improve persuasion, consider Example 1 (respectively Example 2).

*Example 1.* Consider a student health advisor who wants to persuade a student to join a smoking cessation programme (i.e., a health programme designed to help someone give up smoking). The student may be expressing reluctance to join but not explaining why. Through experience, the advisor might guess that the student believes one of the following arguments.

- Option 1: If I give up smoking, I will get more anxious about my studies, I will eat less, and I will lose too much weight.
- Option 2: If I give up smoking, I will start to eat more as a displacement activity while I study, and I will get anxious as I will put on too much weight.

Based on the conversation so far, the student health advisor has to judge whether the student believes option 1 or option 2. With that prediction, the advisor can try to present an appropriate argument to counter the student’s belief in the argument, and thereby overcome the student’s barrier to joining the smoking cessation programme. For instance, if the advisor thinks it is argument 1, the advisor can suggest that as part of the smoking cessation programme, the student can join free yoga classes to overcome any stress that the student might feel from nicotine withdrawal symptoms.

*Example 2.* Consider a doctor in a university health clinic who is trying to persuade a university student to take up regular exercise, and suppose the student says that she does not want to take up a sport because she finds sports boring. The doctor then needs to find a counterargument to the student’s argument. Suppose the doctor has two options:

- Option 1: Doing sport will not only help your physical health, but it will help you study better.
- Option 2: Doing sport will get you in shape, and also help you make new friends.

The argument for Option 1 concerns physical health and getting a good degree, whereas the argument for Option 2 concerns physical health and social life. Now suppose the doctor has learnt through the conversation that the student does not prioritize physical health at all, ranks social life somewhat highly, and ranks getting a good degree very highly. In this case, the doctor will regard the argument in Option 1 as being a better counterargument to present to the student, since it appears to have a better chance of convincing the student.

So in Example 1, the student has the same concerns, but different beliefs, associated with the arguments, whereas in Example 2, the student has the same beliefs, but different concerns, associated with the arguments. We therefore see concerns and beliefs as being orthogonal kinds of information that an agent might have about an argument.

We can use crowdsourcing for the acquisition of user models based on concerns [28] and beliefs [36]. To represent and reason with beliefs in arguments, we can use the epistemic approach to probabilistic argumentation [56, 32, 4, 40, 48] which has been supported by experiments with participants [47]. In applying the epistemic approach to user modelling, we have developed methods for: (1) updating beliefs during a dialogue [33, 34, 39]; (2) efficiently representing and reasoning with a probabilistic user model [25]; and (3) modelling uncertainty in the modelling of persuadee beliefs [35, 27].

### 2.3 Dialogue engine

A **dialogue** is a sequence of **moves** such as asking a query, making a claim, presenting premises, conceding to a premise presented by another agent, etc. The **protocol** specifies the moves that are allowed or required by each participant at each step of a dialogue. There are a number of proposals for dialogues (e.g., [49, 50, 17, 11]). For examples of protocols for persuasion in behaviour change, see [33, 34]. The dialogue may involve steps where the system finds out more about the user’s beliefs, intentions and desires, and where the system offers arguments with the aim of changing the user’s beliefs, intentions and desires. Moves can involve arguments taken from the domain model, and/or they can be queries to improve the user model. In our evaluations (which we review in Section 3), we have focused on the system being able to posit arguments, and the user being able to select his/her counterarguments from a menu of options.

In order to optimize a dialogue (i.e. to maximize the probability that the persuasion is successful), the **strategy** chooses the best moves for the persuader to make in response to the moves made by the persuadee. The strategy model consults the user model to select the moves that are allowed by the protocol. There are a number of roles for arguments. For instance, an argument may be a persuasion goal (i.e., an argument that the system wants the user to accept), or a user counterargument (i.e., an argument that the user regards as a counterargument against an argument by the system), or a system counterargument (i.e., an argument that the system regards as a counterargument against an argument by held by the system), or a user supporting argument (i.e., an argument that the user regards as supporting an argument by held the user), or a system supporting argument (i.e., an argument that the system regards as supporting an argument presented by the system).

There are three options for strategies: The **random strategy** which is a non-deterministic choice of move from available moves. It therefore involves no consideration of the user model; The **local strategy** which involves picking the next move from available moves that is maximal according to some measure of quality based on the beliefs and concerns of the user; And the **global strategy** which involves considering all possible dialogues, and picking the dialogue that maximizes a reward function based on the beliefs and concerns of the user.

We illustrate a local strategy in Example 3, and we use this strategy in the evaluation discussed in Section 3.1.

*Example 3.* We can use a local strategy for taking concerns into account. Consider the following user argument:

- Building cycle lanes is too expensive for the city.

Suppose the following are potential counterarguments with concern assignments given in brackets.

1. (CityEconomy) Evidence shows that infrastructures for cyclists favour the local economy generating more taxes for the city to use.
2. (PersonalEconomy) Cycling is cheaper for the citizens than driving or public transportation.

If the following is a ranking over concerns that is predicted to hold for a given user according to the user model,

$$\text{PersonalEconomy} > \text{Time} > \text{Comfort} > \text{Health} > \text{CityEconomy}$$

then counterargument 2 is the best move.

For a global strategy, our approach to making strategic choices of move is to harness decision trees. A **decision tree** represents all the possible combinations of decisions and outcomes of a sequential decision-making problem. In a situation with two agents, and where the agents take turns, a path from the root to any leaf crosses alternately nodes associated with the proponent (called *decision nodes*) and nodes associated with the opponent (called *chance nodes*). In the

**Table 3.** Some results from the cycling in the city study regarding the proportion of participants going from a negative belief (resp. positive) to a positive belief (resp. negative) and the average change on the participants who did change (avg. change on a scale from -5 to +5).

	Strategic system	Baseline system
Negative to positive	6%	6%
Positive to negative	0%	4%
Avg. change	0.88	-0.14

case of dialogical argumentation, where the proponent (respectively opponent) is a persuader (respectively persuadee), a decision tree represents all possible dialogues. Each path is one possible permutation of the moves permitted by the dialogue protocol *i.e.*, one possible complete dialogue between the two agents. An edge in the tree is the decision (*i.e.*, dialogue move) that has to be taken by the corresponding agent.

Once the decision tree is built, we select, in each decision node, an action to perform (*e.g.*, an argument to posit in each state of the debate) from the point of view of the proponent. This association of a node with the action to perform in this node is called a **policy**. The aim is to compute an **optimal policy**. This is the policy that selects the best action to perform in each decision node. For this, we use a decision rule, composed of two parts: one taking account of the values of all children of a decision node and the other taking account of the values of all the children of a chance node. We can harness decision-theoretic decision rules for optimizing the choice of arguments based on the user model [26, 29]. In Section 3.2, we discuss the evaluation of system that used a global strategy based on decision theory.

Alternatives to our approach for selecting moves include using planning systems [9, 10], minimizing the number of moves [2], selecting a move based on what an agent believes the other is aware of [53], predicting the argument an opponent might put forward based on data about the moves made by the opponent in previous dialogues [24], and using machine learning to predict whether a sequence of dialogue moves would be acceptable to a user [31, 54]. See [57] for a review of strategies in multi-agent argumentation.

### 3 Evaluations with participants

In order to evaluate our framework, we undertook a number of studies with participants [36, 47, 29, 15, 28, 30, 16]. In the following we focus on two of these.

#### 3.1 Cycling in the city study

In this study, we investigated the question of commuting by bicycle in the city [28]. We compiled an argument graph with 51 arguments on the topic of com-



muting by bicycle in the city and 8 concerns (Health, Fitness, Comfort, Time, Personal Economy, City Economy, Environment, and Safety). We undertook pre-studies to validate key assumptions: (1) Participants tend to agree on assignment of concerns to arguments; (2) Participants give meaningful preferences over types of concern; And (3) Participants play by their preferences over concerns. We ran an APS on the web with 100 crowdsourced participants to persuade them to commute by cycle. Using the strategy given in Example 3, we obtained a statistically significant improvement in persuasion when compared with a baseline system that did not consider concerns (see Table 3). This study shows that incorporating concerns can help an APS make better choices of move.


### 3.2 University student fees study

In this study, we investigated the question of university student fees in the UK which normally cost over 9K pounds per annum [30]. We had an argument graph with almost 150 arguments on the topic and 10 concerns (Economy, Government finance, Employment, Student finance, Education, Student satisfaction, Student well-being, University management, Commercialization of universities, Fairness, and Society). We crowdsourced assignment of concerns and beliefs to arguments, and preferences over concerns, for the user model from over 400 participants. We compared our APS with a baseline system (that did not access the user model) using 261 crowdsourced participants where for each participant, if they believed the 9K fee should remain (respectively be abolished), we tried to persuade them that it should be abolished (respectively should remain). We obtained a statistically significant increase in belief in a persuasion goal (average +0.15 on a scale from -3 to +3) when compared with the baseline system. By analyzing the dialogues, the difference in performance is attributable to the better choice of moves made by our APS.

## 4 Towards natural language dialogues

In the evaluation of APSs discussed in Section 3, we did not allow users to type their arguments in natural language. Rather, we presented the user with a menu of potential counterarguments to the previous argument by the system, and the user could select those that s/he subscribed to. Whilst using menus has provided a simple interface in our project, it would be better to have a more natural interface. For argumentation, this means having an interface with some natural language understanding capability. Furthermore, within restricted domains, this can be facilitated by some form of chatbot technology.

A chatbot is a software system with limited natural language processing capability [46]. Simple patterns of normal conversation can be used (e.g. pleasantries). A user can give input in natural language, and this is handled using one or more of the following: simple pattern matching (e.g. regular expressions); natural language parsing; and machine learning classifiers. By determining the type of the user input, the chatbot can then select an appropriate statement as a



**Fig. 1.** Extending the framework for computational persuasion with natural language capability. The arcs denote interactions between the components.

reply from the repository, or ask a query, or change tack by switching to another topic. Such an approach was used by Huang and Lin [31] for participating in dialogues with potential customers with the aim of persuading them to offer a higher price for goods.

#### 4.1 First steps for argumentative chatbots

We now briefly discuss how we can harness the approach of chatbots for argumentation using machine learning techniques. For this, we need to acquire arguments on a topic from crowdsourced participants and then cluster them into groups of similar arguments (as done in [15, 16]). For example, we might have a cluster with sentences that include the following.

- “I don’t exercise because I don’t have a lot of free time during the week.”
- “I am busy doing university work, which is my top priority.”
- “Something always comes up which seems to be more important.”
- “I don’t have enough time.”

An alternative to get a cluster is to start with a domain model (i.e. an argument graph containing all the arguments that the system or user might play). For each argument that the user might play, we can crowdsource linguistic variants of that argument. For this, we present the argument, and ask crowdsourced participants to provide alternative phrasing of the statement. In this way, we can obtain a large number of sentences that contain essentially the same argument (assuming each sentence represents an argument), and we can refer to such a set of sentences as a cluster.

For each cluster (whether obtained by clustering sets of crowdsourced arguments or by obtaining linguistic variants of an argument), we can train a classifier (e.g. a classification tree, a naive Bayes classifier, or a support vector machine) using for instance the SciKit Machine Learning Library for Python. Once trained, the classifier can be applied to user input to determine whether that input belongs to the cluster. For example, if the user types “I work long hours, and there is no space in my schedule for fit in exercise”, the classifier might then classify it as being in the above cluster.

We can aim to have a classifier for each type of counterargument that the user might present in a dialogue on a topic. These classifiers can be harnessed by a state-based chatbot. Each state denotes the state of the dialogue from the point of view of the system. In a state, a statement is selected for presentation to the user. In the simplest case, this can be a non-deterministic choice of candidates that are appropriate for that state of the dialogue. Then the transition to the next state depends on the input by the user, and to what cluster this is classified. So each arc to the next state denotes a type of response made by the user.

For example, in the initial state, the chatbot says “Would you like to talk about exercise?”, and there are two subsequent states. For the first the classifier recognizes positive answers such as “yes” and for the second, the classifier recognizes negative answers such as “no”. At a later stage, the chatbot might give the argument “You need to do more exercise, and so you should consider joining your local gym”. For the subsequent states, there might be classifiers to recognize arguments coming from the user such as about time (as in the bullet points above), or about lack of money, or about lack of interest, etc. Once the classifier has recognized the input, an appropriate answer can be provided.

Therefore, by determining the classification of the user argument (i.e. the cluster to which it belongs), the chatbot can select an appropriate counterargument from the repository, or ask a query, or change tack by using another argument to support the persuasion goal. In addition, simple patterns of normal conversation can be used (e.g. pleasantries).

Furthermore, it is straightforward to implement a simple system that trains classifiers, and incorporates them within state models, so as to allow for simple argumentation dialogues to be undertaken with users in natural language on a narrow topic such as given in Table 2.

## 4.2 Next steps for argumentative chatbots

There are various ways that the classifiers described above could be improved. For our investigations, we used the corpus in [15] which contains clusters of arguments on sufficiently different topics. So for instance arguments about not being able to do exercise because of lack of time can often be discriminated from arguments about not being able to do exercise because of lack of money just by using key words. Indeed, the only features we used for the classifiers were keywords, and synonyms for keywords coming from WordNet in the Python NLTK library [8].

To handle more complex discussions would require more sophisticated discrimination of different arguments. For instance, for the sentence “I am a student, and in my spare time, I prefer to earn money rather than go to the gym”, it is likely to be classified as lack to time or as lack of money. Yet, it seems to fall into a third classification. This therefore calls for richer feature sets for training classifiers which in turn calls for use of bigrams or trigrams [43], if there is sufficient data, or the use of natural language processing to identify syntactic or semantic structure in the input. In particular, the identification of negation,

and the clause within the scope of negation, is an important aspect of understanding counterarguments. Another issue is the pronoun resolution both within a sentence and between the sentence and previous sentences. This creates many challenges, but potentially offers much higher quality classification. Obviously, there is a large literature in natural language processing that is potentially relevant to developing more sophisticated feature sets that should be harnessed.

The design of more sophisticated interfaces could be influenced by developments in argument mining, which is concerned with identifying components of arguments (e.g. premises, claims), and relationships between them (e.g. support, attack), within free text (for reviews see [42, 55]), the use of machine learning to predict the convincingness of arguments [23], and the use of textual entailment to select appropriate responses in an argumentation dialogue [59].

The other aspect of developing argumentation chatbots for persuasion is to hook-up the interface to the dialogue engine so that strategic choices of move can be made based on the domain model and the user model (as illustrated in Figure 1). This could then allow for the menu of counterarguments from which the user selects his/her choice to be replaced by the chatbot natural language interface.

## 5 Discussion

In this paper, we have reviewed a framework for computational persuasion based on domain modelling of arguments and counterarguments, user modelling of the beliefs and concerns of persuadees, and optimizing the choice of move in dialogical argumentation by taking into account the user model. We have discussed studies showing that a system based on this framework can outperform a baseline system over a population of participants.

There are various ways that this framework could be further developed including richer domain models (for example using structured arguments), richer user models (for example using epistemic graphs [38, 37]), and for better methods for strategic arguments (for example better definitions for reward functions).

Then there is the need to develop natural language interfaces so that we are not restricted to menu-driven input from the user. In this paper, we have briefly described a simple approach to harnessing chatbots, and we have described various ways that this could be developed by harnessing developments in computational linguistics. The combination of computational models of arguments (as underlying the framework for computational persuasion) as presented here with computational linguistics could offer some exciting research with important impact. As part of such an endeavor, it is important to ensure that there are studies with participants. This can help to verify the developments are consistent with how people do actually enter into argumentation dialogues.

## References

1. Andrews, P., Manandhar, S., De Boni, M.: Argumentative human computer dialogue for automated persuasion. In: Proceedings of the 9th SIGdial Workshop on

- Discourse and Dialogue, pp. 138–147 (2008)
2. Atkinson, K., Bench-Capon, P., Bench-Capon, T.: Value-based argumentation for democratic decision support. In: Proceedings of ICAART'12. pp. 23–32 (2012)
  3. Atkinson, K., Baroni, P., Giacomin, M., Hunter, A., Prakken, H., Reed, C., Simari, G.R., Thimm, M., Villata, S.: Towards artificial argumentation. *AI Magazine* **38**(3), 25–36 (2017)
  4. Baroni, P., Giacomin, M., Vicig, P.: On rationality conditions for epistemic probabilities in abstract argumentation. In: Proceedings of COMMA'14. pp. 121–132 (2014)
  5. Baroni, P., Gabbay, D., Giacomin, M., van der Torre, L. (eds.): Handbook of Formal Argumentation. College Publications (2018)
  6. Besnard, P., Hunter, A.: Elements of Argumentation. MIT Press (2008)
  7. Bickmore, T., Schulman, D., Sidner, C.: Automated interventions for multiple health behaviours using conversational agents. *Patient Education and Counseling* **92**, 142–148 (2013)
  8. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python — Analyzing Text with the Natural Language Toolkit. O'Reilly Media (2009)
  9. Black, E., Coles, A., Bernardini, S.: Automated planning of simple persuasion dialogues. In: Proceedings of CLIMA'14. LNCS, vol. 8624, pp. 87–104. Springer (2014)
  10. Black, E., Coles, A., Hampson, C.: Planning for persuasion. In: Proceedings of AAMAS'17. pp. 933–942 (2017)
  11. Caminada, M., Podlaszewski, M.: Grounded semantics as persuasion dialogue. In: Proceedings of COMMA'12. pp. 478–485. IOS Press (2012)
  12. Cawsey, A., Grasso, F., Jones, R.: A conversational model for health promotion on the world wide web. In: Proceedings of AIMDM'99. LNAI, vol. 1620, pp. 379–388. Springer (1999)
  13. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolarity in argumentation graphs: Towards a better understanding. *International Journal of Approximate Reasoning* **54**(7), 876–899 (2013)
  14. Chalaguine, L., Hadoux, E., Hamilton, F., Hayward, A., Hunter, A., Polberg, S., Potts, H.: Domain modelling in computational persuasion for behaviour change in healthcare. *ArXiv* (2018), arXiv:1802.10054 [cs.AI]
  15. Chalaguine, L., Hamilton, F., Hunter, A., Potts, H.: Argument harvesting using chatbots. In: Proceedings of COMMA'18. pp. 149–160 (2018)
  16. Chalaguine, L., Hamilton, F., Hunter, A., Potts, H.: Impact of argument type and concerns in argumentation with a chatbot. *ArXiv* (2018), arXiv:1905.00646 [cs.AI]
  17. Fan, X., Toni, F.: Assumption-based argumentation dialogues. In: Proceedings of IJCAI'11. pp. 198–203 (2011)
  18. Fogg, B.: Persuasive computers. In: Proceedings of CHI'98. pp. 225–232 (1998)
  19. Grasso, F.: Exciting avocados and dull pears - combining behavioural and argumentative theory for producing effective advice. In: Proceedings of CogSci'98. pp. 436–441 (1998)
  20. Grasso, F.: Rhetorical coding of health promotion dialogues. In: Proceedings of AIME'03. LNCS, vol. 2780, pp. 179–188 (2003)
  21. Grasso, F., Cawsey, A., Jones, R.: Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition. *International Journal of Human-Computer Studies* **53**(6), 1077–1115 (2000)
  22. Guerrero, E., Nieves, J., Lindgren, H.: An activity-centric argumentation framework for assistive technology aimed at improving health. *Argument and Computation* **7**, 5–33 (2016)

23. Habernal, I., Gurevych, I.: Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM. In: Proceedings of ACL'16. pp. 1589–1599 (2016)
24. Hadjinikolis, C., Siantos, Y., Modgil, S., Black, E., McBurney, P.: Opponent modelling in persuasion dialogues. In: Proceedings of IJCAI'15. pp. 164–170 (2013)
25. Hadoux, E., Hunter, A.: Computationally viable handling of beliefs in arguments for persuasion. In: Proceedings of ICTAI'16. pp. 319–326 (2016)
26. Hadoux, E., Hunter, A.: Strategic sequences of arguments for persuasion using decision trees. In: Proceedings of AAI'17. pp. 1128–1134 (2017)
27. Hadoux, E., Hunter, A.: Learning and updating user models for subpopulations in persuasive argumentation using beta distributions. In: Proceedings of AAMAS'18. pp. 1141–1149 (2018)
28. Hadoux, E., Hunter, A.: Comfort or safety? gathering and using the concerns of a participant for better persuasion. *Argument and Computation* (2019), in press
29. Hadoux, E., Hunter, A., Corrége, J.: Strategic dialogical argumentation using multi-criteria decision making with application to epistemic and emotional aspects of arguments. In: Proceedings of Foiks'18. pp. 207–224. LNCS (2018)
30. Hadoux, E., Hunter, A., Polberg, S.: Strategic argumentation dialogues for persuasion: Framework and experiments based on modelling the beliefs and concerns of the persuadee. Tech. rep., University College London (2019)
31. Huang, S., Lin, F.: The design and evaluation of an intelligent sales agent for online persuasion and negotiation. *Electronic Commerce Research and Applications* pp. 285–296 (2007)
32. Hunter, A.: A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* **54**(1), 47–81 (2013)
33. Hunter, A.: Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In: Proceedings of IJCAI'15. pp. 3055–3061 (2015)
34. Hunter, A.: Persuasion dialogues via restricted interfaces using probabilistic argumentation. In: Proceedings of SUM'16. LNCS, vol. 9858, pp. 184–198 (2016)
35. Hunter, A.: Two dimensional uncertainty in persuadee modelling in argumentation. In: Proceedings of ECAI'16. pp. 150–157 (2016)
36. Hunter, A., Polberg, S.: Empirical methods for modelling persuadees in dialogical argumentation. In: Proceedings of ICTAI'17. pp. 382–389 (2017)
37. Hunter, A., Polberg, S., Potyka, S.: Updating belief in arguments in epistemic graphs. In: Proceedings of KR'18). pp. 138–147 (2018)
38. Hunter, A., Polberg, S., Thimm, M.: Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *ArXiv* (2018), arXiv:1802.07489 [cs.AI]
39. Hunter, A., Potyka, N.: Updating probabilistic epistemic states in persuasion dialogues. In: Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'17). LNCS, vol. 10369, pp. 46–56. Springer (2017)
40. Hunter, A., Thimm, M.: Optimization of dialectical outcomes in dialogical argumentation. *International Journal of Approximate Reasoning*, **78**, 73–102 (2016)
41. Hunter, A.: Towards a framework for computational persuasion with applications in behaviour change, argument and computation. *Argument and Computation* **9**(1), 15–40. (2018)
42. Lippi, M., Torroni, P.: Argumentation mining: State of the art and emerging trends. *ACM Transactions on Internet Technology* (2016)
43. Manning, C., Schütze, H.: *Foundations of Statistical Natural Language Processing*. MIT Press (2000)

44. Nguyen, H., Masthoff, J., Edwards, P.: Persuasive effects of embodied conversational agent teams. In: Proceedings of Human-Computer Interaction. LNCS, vol. 4552, pp. 176–185 (2007)
45. Ogden, J.: Health Psychology: A Textbook. Open University Press, 5th edn. (2012)
46. Perez-Marin, D., Pascual-Nieto, I. (eds.): Conversational agents and natural language interaction; techniques and effective practices. Information Science Reference (2011)
47. Polberg, S., Hunter, A.: Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning* **93**, 487–543 (2018)
48. Polberg, S., Hunter, A., Thimm, M.: Belief in attacks in epistemic probabilistic argumentation. In: Proceedings of SUM'17. LNCS, vol. 10564, pp. 223–236 (2017)
49. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* **15**(6), 1009–1040 (2005)
50. Prakken, H.: Formal systems for persuasion dialogue. *Knowledge Engineering Review* **21**(2), 163–188 (2006)
51. Rahwan, I., Madakkatel, M., Bonnefon, J., Awan, R., Abdallah, S.: Behavioural experiments for assessing the abstract argumentation semantics of reinstatement. *Cognitive Science* **34**(8), 1483–1502 (2010)
52. Rahwan, I., Simari, G. (eds.): *Argumentation in Artificial Intelligence*. Springer (2009)
53. Rienstra, T., Thimm, M., Oren, N.: Opponent models with uncertainty for strategic argumentation. In: Proceedings of IJCAI'13. pp. 332–338 (2013)
54. Rosenfeld, A., Kraus, S.: Providing arguments in discussions on the basis of the prediction of human argumentative behavior. *ACM Transactions on Interactive Intelligent Systems* **6**(4), 30:1–30:33 (Dec 2016)
55. Stede, M., Schneider, J.: *Argumentation Mining*. Morgan & Claypool (2019)
56. Thimm, M.: A probabilistic semantics for abstract argumentation. In: Proceedings of ECAI'12. vol. 242, pp. 750–755 (2012)
57. Thimm, M.: Strategic argumentation in multi-agent systems. *Kunstliche Intelligenz* **28**, 159–168 (2014)
58. Thomas, R.J., Masthoff, J., Oren, N.: Adapting healthy eating messages to personality. In: Proceedings of Persuasive'17. LNCS, vol. 10171, pp. 119–132 (2017)
59. Toniuc, D., Groza, A.: Climebot: an argumentative agent for climate change. In: Proceedings of ICCP'17. pp. 63 – 70 (2017)
60. Zaal, E., Mills, G., Hagen, A., Huisman, C., Hoeks, J.: Convincing conversations: Using a computer-based dialogue system to promote a plant-based diet. In: Proceedings of the Cognitive Science Conference (CogSci'17). pp. 3627–3632 (2017)