

Computational Persuasion: Towards Applications in Behaviour Change

Anthony Hunter¹

¹Department of Computer Science, University College London, UK.

July 1, 2016

Questions that we will consider

- What is persuasion?
- What is computational persuasion?
- What do computational models of argument offer?
- How can computational persuasion be developed for applications in behaviour change?



Computational models of argument (peg) and computational persuasion (hole)

What is persuasion?

What is persuasion?

Process by which one agent intends to induce another agent to undertake a particular physical or mental action.

Some examples of unidirectional persuasion

- Product advertising (e.g. T.V. or magazine advert)
- Political speech (e.g. prior to an election)
- Government advisory messages (e.g. recycle paper, metal, etc).

20,679* Physicians say "LUCKIES are less irritating"

"It's toasted"

Your Throat Protection against irritation against cough

*The figures quoted herein are based on research conducted by the American Medical Association and the American Lung Association.

What is persuasion?

What is persuasion?

Process by which one agent intends to induce another agent to undertake a particular physical or mental action.

Some examples of bidirectional persuasion

- Sales meeting (e.g. in car showroom)
- Some kinds of medical counselling (e.g. drug abuse)
- Discussions with a goal (e.g. employee asking for a payrise)



7121 [RF] © www.visualphotos.com

What is persuasion? — Rules of the game

Rule 1: The only goal is to change to mind of persuadee

Rule 2: There are no further rules

Some observations concerning these rules

- Argumentation does not need to be normative (so arguments can be inconsistent, irrational, untrue, etc).
- Though inconsistent, irrational, untrue arguments may be counter-productive with some audiences.


What is persuasion? — Persuader can be important

Seemingly good features

- Authority
- Expert
- Knowledge

Seemingly poor features

- Attractive
- Witty
- Celebrity



What is persuasion? — Persuader can be important



Safety gear for bikes

- Teenager is unlikely to be convinced by a seemingly boring government expert on bike safety
- Teenager is more likely to be convinced by someone like a sports star, or celebrity.



What is persuasion? — Language can be important

Some linguistics dimensions of persuasion

- Choice of words and phrases (e.g. “field sports” vs “blood sports”, or “freedom fighter” vs “terrorists”).
- Metaphor (e.g. using “the whole world’s a stage” when persuading someone to do something bold).
- Metonymy (e.g. Oscar Wilde on fox hunting “The unspeakable in pursuit of the uneatable”).
- Irony (e.g. Denis Healey describing an attack by Geoffrey Howe as “like being savaged by a dead sheep”).



Cockcroft & Cockcroft (1992) *Persuading People: An Introduction to Rhetoric*

What is persuasion? — Psychology can be important

Some strategies based on empirical evidence

- 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)
- Social proof (e.g. treating dog phobia in children by showing videos of children playing happily with children)




Cialdini (1997) *The Psychology of Influence*

What is persuasion? — Personality can be important

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.
- 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.



ARGUMENTATION
IS AT THE HEART
OF PERSUASION



Some kinds of interaction surrounding persuasion

- 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 arguments (and counterarguments) are the essential structures for presenting the claims (and counter claims) in persuasion**

What is persuasion? — Being rational can be important

Some criteria for the persuadee being convinced by a persuasion argument

- Acceptability of persuasion argument (against counterarguments)
- Believing the premises of the persuasion argument
- Fit of persuasion argument with agenda, goals, preferences, etc
- Quality of arguments (balance, depth, breadth, understandable, etc)



What is persuasion? — Need for convincing arguments

How convincing an argument is \neq How correct it is.

- Homeopathy focuses on processes of health and illness rather than states, and therefore it is better than regular medicine
- 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



What is persuasion? — Arguments can be emotional

Emotional arguments play on the emotions of the persuadee

- 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.
- Your parents will be proud of you if you complete your thesis and get your PhD award.



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*).

What is persuasion? — Selectivity is important

Example where initial argument is not believed by other person

- **Him** “The car is a nice red colour, and that is the only criterion to consider, therefore we should buy it.”
- **Her** “It is a nice red colour, but I don’t agree that that is the only criterion to consider.”

Example where initial argument is believed by other person

- **Him** “The car is the most economical and easy car to drive out of the options available to us, and those are the criteria we want to satisfy, so we should buy the car.”



What is persuasion? — Conclusions

Success may depend on

- an appropriate persuader
- an appropriate language
- personality of persuadee
- use of psychological techniques
- use of rational argumentation
- use of emotional arguments
- selectivity in arguments presented based on model of persuadee
 - bias/personality
 - beliefs
 - awareness of arguments
 - preferences/agenda



Computational models of argument focus mainly on use of rational argumentation.


What is computational persuasion?

Definition of automated persuasion system

An **automated persuasion system** (APS) is a system that can engage in a dialogue with a persuadee in order to persuade the persuadee to do (or not do) some action or to believe (or not believe) something via argumentation.


Definition of computational persuasion

Computational persuasion is the study of formal models of dialogues involving arguments and counterarguments, of user models, strategies, etc. for APSs.



What do computational models of argument offer?

Levels of complexity in computational models



What do computational models of argument offer?

Components of a model of dialogical argumentation

- Participants** Specification of the information held by each agent (e.g. a knowledgebase, a set of goals, etc.)
- Moves** Specification of the moves that can be made (e.g. `why(ϕ)`, `claim(ψ)`, `posit(A)`, etc.)
- Protocol** The rules of the game (i.e. the moves an agent is allowed, or is obliged, to make at each stage of the dialogue).

See Hamblin (Theoria 1971); MacKenzie (JPL 1979)


What do computational models of argument offer?

Dialogue between Paul and Olga about Paul's claim that a particular car is safe.

- 1 `claim(Paul, safe)`
- 2 `why(Olga, safe)`
- 3 `explain(Paul, {airbag, airbag \rightarrow safe}, safe)`
- 4 `concede(Olga, airbag)`
- 5 `explain(Olga, {explosions, explosions \rightarrow \neg safe}, \neg safe)`
- 6 `explain(Paul, {unreliablenews,
unreliablenews \rightarrow \neg explosions}, \neg explosions)`
- 7 `explain(Olga, {highspeed, highspeed \rightarrow \neg safe}, \neg safe)`

Example adapted from Prakken (KER 2006)

What do computational models of argument offer?



What do computational models of argument offer?

Argument-based persuasion in the literature

- Protocols for logic-based argumentation in persuasion dialogues (e.g. [Angoud, Parsons & Maudet 2001](#), [Prakken 2006](#), [Fan & Toni 2011](#), etc) but assume exhaustive presentation of counterarguments.
- Persuadee modelled by uncertainty over structure of argument graph ([Nir, Atkinson & Li, 2012](#)) but no consideration of dialogue or strategy.
- Persuadee modelled by uncertainty over what s/he knows about, and this is used to choose next move ([Rienstra, Thimm & Oren, 2013](#)) but no consideration of beliefs.
- Persuadee modelled by uncertainty over arguments likely to be presented based on previous dialogues ([Hadjinikolis et al, 2013](#)) but no consideration of belief.

Promising ideas on strategies - but as we will see, we need to take belief into account, we need viable algorithms, we need to take more features of the audience into account, etc., etc.

What do computational models of argument offer?

Some proposals for taking the audience into account

- Value-based argumentation in which the ethical value of each argument is taken into account (Bench-Capon 2003)
- Taking beliefs audience into account when assessing the empathy/antipathy that audience has in each argument (Hunter 2004).
 - E.g. Consider a politician justifying tax rise to different audiences such as business people, young people, old people, etc.
- Taking goals of audience into account when generating threat or reward arguments (Amgoud 2005).
 - E.g. Persuading people to recycle by giving a rebate on their local tax.

No consideration of how such methods can be harnessed in a persuasion strategy.



What do computational models of argument offer?

Argument dynamics (changing graph to ensure specific outcomes)

- Epistemic enforcement in abstract argumentation (e.g. Baumann & Brewka 2012, Coste-Marquis et al 2014/2015)
- Belief revision in abstract argumentation (e.g. Cayrol et al 2010, Gabbay & Rodrigues 2012, Bisquert 2013, Diller et al 2015)

Example

Given the left graph, suppose we want A_2 and A_4 in the grounded extension.



No consideration of how such methods can be harnessed in a persuasion strategy.

What do computational models of argument offer?

Many important developments in abstract, logical, and dialogical argumentation, and in argument dynamics.


Studies with participants (e.g. Rahwan et al, Cerutti et al, Rosenfeld & Kraus) are promising for grounding these developments.

How do we move forward?

- Need further development of opponent modelling, strategies, etc.
- Need to address the interface problem (i.e. overcome the natural language problem).
- Need to test our ideas.




Computational persuasion in behaviour change



- Persuasion technologies are being developed to help people make positive changes to their behaviour (e.g. healthcare and healthy life styles).
- There is an emphasis on either helping users to explore their issues (e.g. game playing) or helping users once they are persuaded to do something (e.g. diaries for recording calorie intake for weight management).
- Interestingly, argumentation is not central to the current manifestations of persuasion technologies.

Some potential domains for argument-based persuasion technology

- healthy life styles (e.g. eating fewer calories, eating more fruit and veg, exercise, drinking less alcohol)
- addiction management (e.g. gambling, smoking, drugs)
- treatment compliance (e.g. self-management of diabetes, taking vaccines, completing course of antibiotics)
- personal finance (e.g. borrowing less, saving more)
- education (e.g. starting or continuing with a course, studying properly)
- energy efficiency (e.g. reducing domestic electricity consumption, installing home insulation)
- citizenship (e.g. voting, recycling, giving to charities, wasting less food)
- safe driving (e.g. not exceeding speed limits, not texting while driving)
- anti-social behaviour (e.g. aggression, vandalism, racism, sexism, trolling)



Our requirements for computational persuasion via an app

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

Stages of Change model

Pre-contemplation “I am happy being a smoker and intend to continue smoking”

Contemplation “I have been coughing a lot recently, perhaps I should think about stopping smoking”

Preparation “I will buy lower tar cigarettes”

Action “I have stopped smoking”

Maintenance “I have stopped smoking for four months now”

Prochaska & Diclemente (1982)

Example from Ogden (2011) *Health Psychology*

Computational persuasion in behaviour change



- 1 (System) You believe that a cup cake is preferable to a banana?
- 2 (User) Yes.
- 3 (System) It is late afternoon, and you think a cup cake will give you a sugar rush to help you work?
- 4 (User) Yes.
- 5 (System) The sugar rush from a cup cake is brief, and therefore it won't help you work.
- 6 (System) A banana gives a longer lasting energy supply, and so a banana is preferable to a cup cake.

Computational persuasion in behaviour change

Natural language dialogues are not possible in the short-term.

Some options for asymmetric dialogues

- Persuader can posit arguments.
- Persuader can ask questions and the persuadee answers (e.g. yes/no).
- Persuadee is presented counterarguments and has to agree/disagree.

Since you do little exercise, you should do a regular exercise class

When I do exercise, I get very hungry and I put on weight

Strongly agree

Agree

Neither agree nor disagree

Disagree

Strongly disagree

Computational persuasion in behaviour change

How often do you have face-to-face contact with patients?

Less than once per week

Between 1 and 5 per week

Between 6 and 30 per week

Between 30 and
100 per week

More than 100 per week

Which are goals for you?

Staying healthy

Living a long life

Doing interesting activities

Relaxing

Doing as little as possible

Being with the family

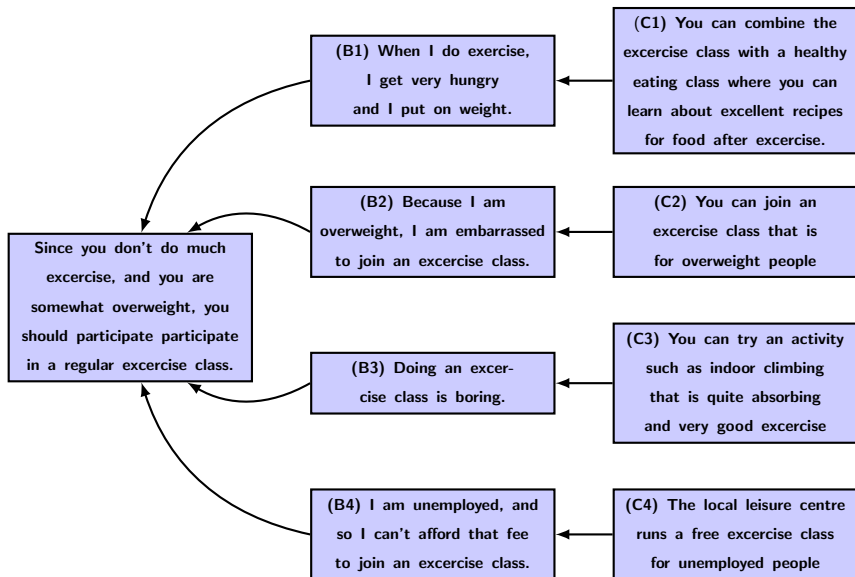
Being healthy for the family

Looking good

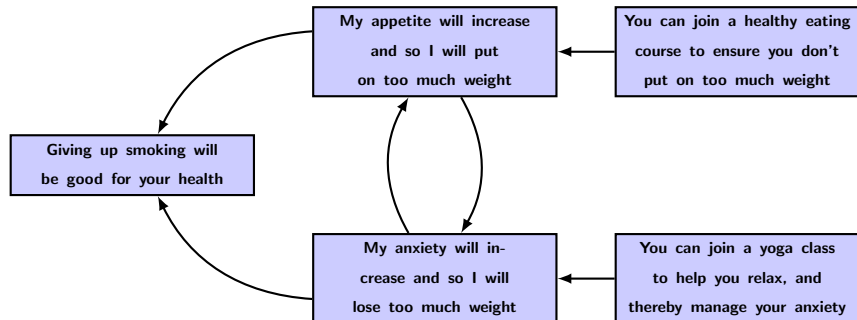
Being fit for my sport

Saving time

Computational persuasion in behaviour change



Computational persuasion in behaviour change



Computational persuasion in behaviour change

Diverse kinds of counterarguments to be addressed

- 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. driving on a race track)
- 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)
- Etc.



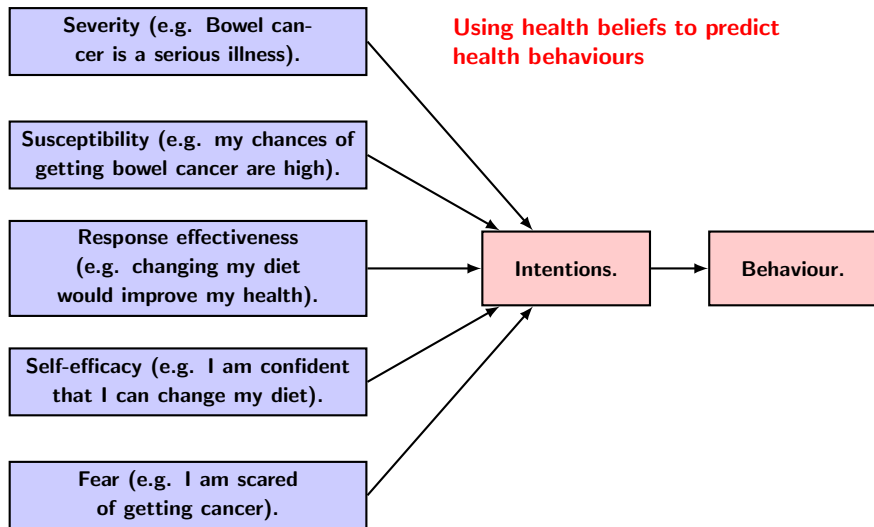
We aim to design an ontology of argument types to use in protocols and strategies.

Arguments are automatically generated from a knowledgebase.

- There are many commonalities in the knowledge required for each behavioural change application.
 - Persuadee beliefs (e.g. cakes give a sugar rush)
 - Persuadee preferences (e.g. apples are preferred to oranges, burgers are preferred to apples)
 - Behavioural states (e.g. persuadee's weight, persuadee's typical exercise regime)
 - 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 refined sugar intake by 90%).

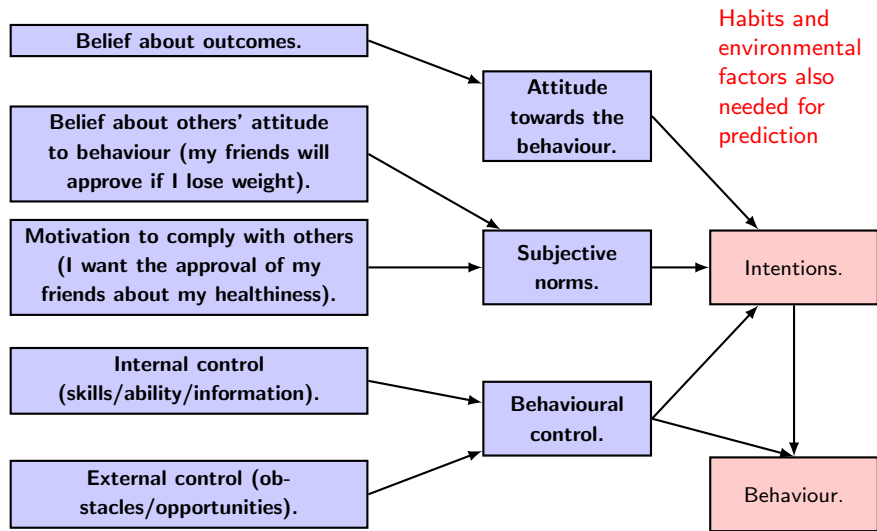
We aim to develop a calculus in predicate logic (drawing on the BDI approach) for relating beliefs, behavioural goals, and behavioural states, to possible actions.

Computational persuasion in behaviour change



Protection motivation theory (Roger 1975)
Example from Ogden (2011) *Health Psychology*

Computational persuasion in behaviour change



Theory of planned behaviour (Ajzen 1991)
Example from Ogden (2011) *Health Psychology*

Dimensions of uncertainty:

- Arguments/attacks known by persuadee
- Beliefs of persuadee
- Moves that persuadee might make
- Risk of disengagement



Computational persuasion in behaviour change

Constellations approach

Let G be an argument graph, let \sqsubseteq be the subgraph relation, and let P be a probability distribution.

$$P : \{G' \sqsubseteq G\} \rightarrow [0, 1]$$

	Subgraph	Probability
G_1	$A \leftrightarrow B$	0.09
G_2	A	0.81
G_3	B	0.01
G_4		0.09

$$\begin{aligned} P_{\text{gr}}(\{A, B\}) &= &= 0.00 \\ P_{\text{gr}}(\{A\}) &= P(G_2) &= 0.81 \\ P_{\text{gr}}(\{B\}) &= P(G_3) &= 0.01 \\ P_{\text{gr}}(\{\}) &= P(G_1) + P(G_4) &= 0.18 \end{aligned}$$

Use in asymmetric persuasion

- Constellation approach can model the uncertainty about the structure of the graph in the persuadee mind.
- Update the model with each argument/attack presented.
- Use expected utility to identify best choice of argument/attack to present (see [Hunter & Thimm COMMA'14](#), [Int J. Approx Reasoning 2016](#)).

Computational persuasion in behaviour change

Epistemic approach

Let P be a probability distribution.

$$P : \wp(\mathcal{A}) \rightarrow [0, 1]$$

Example of the epistemic approach

Suppose I hear one of my friends saying argument A and another saying argument B .

A = John suffers from hay fever, and so a picnic in the hay field will be unpleasant for him.

B = John has taken a homeopathic medicine for hay fever and therefore he won't suffer from hay fever.



If I believe that homeopathic medicine is just water, then I have high belief in A and low belief in B (e.g. $P(A) = 0.9$ and $P(B) = 0$).

Computational persuasion in behaviour change

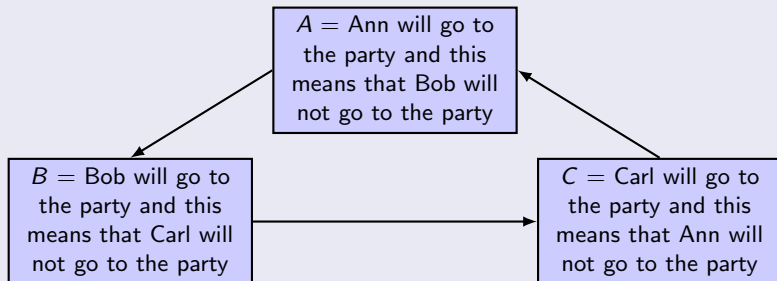
Definition

For an argument graph G , and a probability assignment P , the **epistemic extension** is

$$\{A \in \mathcal{A} \mid P(A) > 0.5\}$$

Example

Suppose we have $P(A) = 0.9$, $P(B) = 0.1$, and $P(C) = 0.1$, then the epistemic extension is $\{A\}$.




Computational persuasion in behaviour change

Definition

A probability function P is **rational** for an argument graph $(\mathcal{A}, \mathcal{R})$ iff for each $(A, B) \in \mathcal{R}$, if $P(A) > 0.5$, then $P(B) \leq 0.5$.

Example



Some examples of probability functions.

A	B	C	rational?	epistemic extension
0.3	0.1	0.9	yes	$\{C\}$
0.9	0.1	0.9	yes	$\{A, C\}$
0.1	0.8	0.1	yes	$\{B\}$
0.1	0.8	0.9	no	$\{B, C\}$
0.7	0.8	0.5	no	$\{A, B\}$

[See Hunter 2013]

The epistemic approach can give a finer grained version of Dung's approach to obtaining extensions, and it can be used to give a natural alternative to Dung's extensions (see Hunter & Thimm ECAI'14).

Using epistemic approach to model beliefs of persuadee (Hunter IJCAI'15)


- The epistemic approach is useful for asymmetric dialogues where the user is not allowed to posit arguments or counterarguments.
 - 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).

[See Hunter IJCAI'15, ECAI'16, SUM'16].

Computational persuasion in behaviour change

Schematic representation of a dialogue and the user models

Let $D = [m_1, \dots, m_n]$ be a dialogue



Each user model P_i is obtained from P_{i-1} and m_i using an update method.

Aim of dialogue w.r.t. persuasion goal ϕ (a Boolean combination of arguments)

Maximize $P_n(\phi)$ (i.e. according to user model, the user believes ϕ)

Computational persuasion in behaviour change

The strict method for updating the user model

For step i in the dialogue, the **strict method** generates P_i from P_{i-1} as follows, where $\Phi = \{\neg C \mid (A, C) \in \text{Attacks}(G)\}$.

If $D(i) = A!$,
and for all $(B, A) \in \text{Attacks}(G)$, $P_{i-1}(B) \leq 0.5$,
then $P_i = \text{Update}_{\Phi}^1(\text{Update}_A^1(P_{i-1}))$,

Example




For dialogue $[A!, C!, A!]$, the trusting method gives the following updates.

	111	110	101	100	011	010	001	000
P_0	0.2	0.3			0.3	0.2		
P_1	0.2	0.3			0.3	0.2		
P_2			0.5				0.5	
P_3			1.0					

Computational persuasion in behaviour change

We can model the moves that an agent might make using a (probabilistic) finite state machine with each state being a tuple. This can be explored using MINIMAX (see [Hunter SUM'13,SUM'14](#)).

(PersuaderPrivateState, PublicState, PersuadeePrivateState)



We can find optimal sequences of moves by handling uncertainty concerning the persuadee using partially observable markov decision processes (POMDPs) (see [Hadoux et al IJCAI'15](#)).

Computational persuasion in behaviour change

Dimensions of uncertainty in models of persuadee:

- Beliefs of persuadee (**Epistemic approach**)
- Arguments/attacks known by persuadee (**Constellations approach**)
- Moves that persuadee makes (**PFSMs/POMDPs**)
- Risk of disengagement (**Markov model**)

Need for a deeper understanding of the relationships between these dimensions.

Ongoing development of strategies based on:

- Domain knowledge for constructing rational and emotional arguments
- Protocols for asymmetric dialogues
- Models of persuadee

Computational persuasion in behaviour change

We aim to evaluate prototype systems (% of users persuaded) in two domains.

Weight + exercise behaviour

- Over 60% of adults in England are overweight or obese.
- Overweight people have higher risk of type 2 diabetes, heart disease and certain cancers.
- Health problems associated with being overweight or obese cost the NHS more than 5 billion every year.

Flu vaccination behaviour

- NHS England has 1.3 million employees.
- Flu causes many lost days of work.
- NHS employees are vulnerable to infection from patients with flu.
- Flu passed onto patients causes complications.
- Only 50% of NHS employees take flu vaccine.

Domain knowledge for the arguments comes from medical literature on trials, clinical guidelines, clinical protocols, and data from consultations.

Conclusions

Computational persuasion based on computational models of argument is a promising approach to technology for behavioural change applications.

Ongoing work

- Richer models of the user.
- Methods for learning user models from data.
- Richer strategies.
- Richer asymmetric dialogues.
- Applications in behaviour change.



- This talk is based on
 - A. Hunter (2016) Computational Persuasion with Applications in Behaviour Change, in Computational Models of Argument, IOS Press, (in press).
- Further papers include
 - A. Hunter (2014) Opportunities for Argument-Centric Persuasion in Behaviour Change, JELIA 2014: 48-61.
 - A. Hunter (2015) Modelling the Persuadee in Asymmetric Argumentation Dialogues for Persuasion, IJCAI 2015: 3055-3061.
 - A. Hunter (2016) Two Dimensional Uncertainty in Persuadee Modelling in Argumentation, ECAI 2016. (in press).
 - A. Hunter (2016) Persuasion Dialogues via Restricted Interfaces using Probabilistic Argumentation, SUM 2016. (in press).