

From Psychological Persuasion To Abstract Argumentation: A Step Forward

Emmanuel Hadoux¹, Ariel Rosenfeld² and Jean-Baptiste Corrége³

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

² Weizmann Institute of Science, Rehovot, Israel

³ LIMSI, CNRS, Université Paris-Saclay, F-91405 Orsay

Abstract. Developing argumentation-based persuasive agents that leverage human argumentative techniques is an open challenge in the computational argumentation field. In this paper, we propose a computational perspective on the psychological techniques people tend to follow during persuasion interactions drawing on psychological evidence. We focus on four well-established psychological techniques, model and investigate them using a recently proposed argumentative computational framework. Our investigation reveals both similarities and gaps between the two which can be either leveraged or addressed in the design of argumentation-based persuasive agents and future theoretical developments.

1 Introduction

A key human skill, used across many domains and activities, is the ability to *persuade*. Politicians strive to persuade their constituents, colleagues try to persuade each other to adopt different solution techniques for a joint problem, parents try to persuade their children to eat healthier food, etc. People use many different techniques for persuading others. These *human persuasive techniques* have been thoroughly investigated in the *real world* by psychology researchers. Surprisingly, despite the major advancements of the computational **argumentation theory**, providing grounded techniques and models analyzed and tested in theoretical settings, the study of the possible connections between human persuasive techniques and computational models has yet to be properly examined.

In this work we provide a novel investigation of the connections between psychological persuasion literature and argumentation theory. Through this investigation we are able to identify the potential use of psychological persuasive principles in argumentation-based systems and find potential directions for future work in adapting and/or extending current argumentative principles to correctly account for psychological persuasion literature. *Our findings contribute an additional stage in the greater challenge of bridging the gap between argumentation theory and people.*

The paper is structured as follows: In Section 2 we survey related works which tried to bridge the gap between psychology and argumentation, coming from both sides. We also review the *Weighted Attack/Support*

Argumentation graphs [MN16] and the necessary definitions used in this paper. In Section 3, we discuss four well-established psychological persuasive techniques. For each technique, we present the idea underlying the technique along with supportive evidences validating the technique. Finally, in Section 4, we model the psychological persuasive techniques discussed in Section 3 using abstract argumentation and evaluate the resulting model.

2 Background

Within the computational argumentation field, a significant effort has been placed on proposing and evaluating models and techniques aimed at allowing an automated agent (*i.e.*, *persuader*) to persuade a person (*i.e.*, *persuadee*). Theoretically, an agent would seek to deploy an *optimal* persuasive policy, mapping each possible state of a dialogue to the *best* argument for the agent to present. This persuasive policy may strive to maximize different objectives:

- likelihood of having a specific set of arguments (*i.e.*, goal arguments) accepted at the end of the dialogue, [HBM⁺15,BCH17].
- persuadee’s valuation of a specific point of view (represented as a single goal argument) [RK16b].
- belief of the persuadee in the goal arguments [HH17],
- plausibility of the goal arguments [MN16].

However, while different computational argumentative techniques have been proposed and investigated in theoretical settings, *human persuasive techniques* have been thoroughly investigated in the *real world* by psychology researchers. These studies have identified the psychological grounds and characteristics of the different techniques that people actually use. The apparent gap between the notion of persuasion in argumentation theory and human persuasive techniques prevents automated persuasive agents from building upon proven psychological persuasive evidence and thus reduces the potential impact of such agents.

A handful of previous works have examined different facets of the connections between argumentation theory and human behaviour. For example, Rahwan et al. [RMB⁺10] have studied the reinstatement argumentative principle in questionnaire-based experiments, Cerutti et al. [CTO14] examined humans’ ability to comprehend formal arguments and Rosenfeld and Kraus [RK14,RK16a] have established that the argumentation theory falls short in explaining people’s choice of arguments in synthetic and real world argumentative settings. To the best of our knowledge, in this recent line of research, no work has used *psychological evidences* to investigate the computational argumentation theory applicability and the possible adaptation thereof.

In order to perform reasoning in a persuasive context, an argumentation framework needs to be defined (see [BPW14] for a recent review). In its most basic form, an argumentation framework consists of a set of arguments A and an attack relation R over $A \times A$ [Dun95]. In previous




Fig. 1: Example of bipolar argument graph where plain arrows mean attacks and dashed arrows mean supports.

investigations of human argumentative behaviour (*e.g.*, [RK14,RK16a]), it was noticed that people often use supportive arguments rather than attacking ones, which necessitates the addition of the support relation as suggested in [ACLSL08]. Furthermore, it is shown that people associate different belief levels in arguments, as suggested in [BRT⁺15], and different strength levels with interactions between arguments, as suggested in [DHM⁺11]. Interestingly, a framework named *Weighted Attacks/Support Argument* [MN16] embedding all these components has recently been proposed. We review this framework below.

Weighted Attacks/Support Argument

Weighted Attacks/Support Argument (WASA) graphs [MN16] are able to model argument graphs with attacks, supports, initial plausibility and strength of interactions between arguments taken into account. This framework merges several concepts: First of all, it is bipolar [ACLSL08], allowing an additional support relation. Moreover, it uses initial weights as the plausibility for the arguments. In this work, we interpret the plausibility as an initial strength given to an argument.

A WASA graph is characterized by a triplet $\mathbb{A} = \langle \mathcal{A}, \mathcal{G}, w \rangle$, where,

- \mathcal{A} is a vector of size n ordering a set of arguments,
- \mathcal{G} , the *incidence matrix*, a square matrix of order n , with $g_{ij} \in \{-1, 0, 1\}, \forall i, j \in \{1, \dots, n\}$, where $g_{ij} = -1$ (resp. 1) represents an attack (resp. a support) from j to i and 0 means no relation,
- w is a weight vector in \mathbb{R}^n .

Example 1. Example of WASA graph.

The bipolar argument graph depicted in Figure 1 can be represented as a WASA graph as follows:

$$\left\langle \begin{pmatrix} a \\ b \\ c \end{pmatrix}, \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} w_0^a \\ w_0^b \\ w_0^c \end{pmatrix} \right\rangle$$

The acceptability of an argument a is called the *acceptability degree* $Deg_{\mathbb{A}}(a)$ in a WASA \mathbb{A} .

The calculation of the acceptability degree in a WASA $\mathbb{A} = \langle \mathcal{A}, \mathcal{G}, w \rangle$ is as follows. First we need to define a *damping factor* $d \geq 1$, acting as a decreasing effect the further the arguments are from a in the argument graph. Then, we can calculate the propagation matrix

$$P_r^{\mathcal{G},d} = \sum_{i=0}^{\infty} \left(\frac{1}{d} \mathcal{G} \right)^i$$

with \mathcal{G} the incidence matrix. Note that the sum is defined at the infinite. However, it is enough to sum until it converges to a stable propagation matrix. This is the case if $d > indegree(\mathcal{G})$.

Finally, we can calculate the acceptability degree vector

$$Deg_{\mathbb{A},d} = Pr^{\mathcal{G},d} \times w$$

for all the arguments.

In this work, we extend the traditional argument graph depiction as presented in Figure 1 to take into account temporal aspects of the dialogue. Namely, we add additional information to the graph: the *step* at which the argument has been/has to be played. This also allows us to represent duplicate arguments that may be played several times in a given dialogue. Specifically, each argument is amended with a subscript, denoting the order in which arguments are presented. A subscript of zero denotes that an argument is not presented at all.

3 Psychosocial Persuasion Principles

We focus on four well-established techniques commonly used by professional in, for instance, sales or marketing, which have been formalized by psychologists. These technique are aimed at persuading other people. Within the field of persuasive technologies, Fogg [Fog98] defined persuasion as “an attempt to shape, reinforce, or change behaviours, feelings, or thoughts about an issue, object, or action”.

Each of the four techniques is presented along with the psychological intuition standing behind it and one or two human studies from the literature that corroborated the benefit of the technique. Following psychological terminology, we define the *goal request* to be a designated argument in the argumentation framework which represents the persuader’s aim or goal – namely, having the persuadee doing or believing something. This goal request is equivalent to the goal argument in Rosenfeld and Kraus’s framework [RK16b]). Positing only the goal argument would probably not suffice to persuade the persuadee in many setting. Therefore, it is necessary to posit additional arguments which interact with the goal request or interact with other arguments that may attack the goal request. Each technique prescribes a procedure of how and when to posit these additional arguments.

3.1 Foot in the Door

The Premise The *foot in the door* technique has been first described by Freedman and Fraser [FF66]. Using this technique, a small initial request serves as a preparatory action for the goal request. Individuals who have been asked a small request generally tend to answer more favourably compared to individuals who have been asked the goal request. This effect is due to the fact that accepting a small, initial request leads individuals to see themselves as being social – “agreeing to requests made by strangers”. Consequently, when confronted with a second request, individuals tend to comply with the above perception and accept more

willingly a bigger request. Failure to conform to the self-image generated by the first request generates a cognitive dissonance, which can explain compliance.

Studies In their paper, Freedman and Fraser report two studies. In the first one, the goal request was to ask housewives to allow a survey team to come into their homes for 2 hours in order to conduct a study about the household products they use. Participants were divided in four categories, depending on the first contact (i.e., the initial request) before asking the goal request:

1. they were asked to answer some short questions about the kinds of soaps they use.
2. they were asked if they would be willing to answer different questions but the questions themselves were not asked.
3. they were merely approached but not asked anything.
4. there was no initial contact (control group).

Results show that the compliance rate with the survey request is:

1. 52.8% when the *foot in the door* was used,
2. 33.3% when firstly engaged,
3. 27.8% when merely approached,
4. 22.2% for the control group.

In their second study, the goal request was to ask participants to put a very large sign which said “Drive Carefully” in their front yard. The authors designed several types of initial requests (e.g., participants were initially asked to either put a small sign in their garden or sign a petition). A control condition was added, in which participants were not initially approached. In the control condition, only 16.7% of the participants complied with the goal request. The highest compliance rate was obtained by asking something that was similar (i.e., put a small sign) and on a similar issue (i.e., safe driving), in which case 76% of the participants agreed to the goal request. In the three other configurations, 47% complied with the goal request, which remains higher than in the control condition.

These results show that making a small initial request before a larger one brings about an increased compliance rate with the goal request. This effect holds whether both requests focus on the same behaviour or not and whether both requests target the same issue or not. However, the best compliance rate is achieved when both requests target the same type of behaviour, focused on the same issue.

3.2 Door in the Face

The Premise The *door in the face* principle has been first theorized by Cialdini et al. [CVL⁺75]. This technique is the almost symmetrical to the *foot in the door* technique discussed above. Using the *door in the face* technique one asks an “unreasonable” request before proposing a smaller one – which is the goal request. The mechanism behind this technique is that, after the big request have been rejected, proposing a

smaller request is perceived as a concession that the persuader has made from her original request. Thus, in order to maintain a certain level of reciprocity in the relation, the persuadee will tend to comply more with the target request than if it was made without the preparatory action.

Studies In their original paper [CVL⁺75], the authors report three studies. In one of these studies, the goal request was to have students accompany a group of juvenile delinquents on a two-hour trip to the zoo. They grouped the students into three conditions:

1. students with whom they engaged the interaction by asking them first to act as counselors to juvenile delinquents for a period of two years,
2. students without any other request except for the two-hour trip (control group),
3. students where both options were presented and subjects were to choose which (if any) of the two options to take.

The results are as follows:

1. 50% agreed when the door in the face technique was used (first asking the big request).
2. 16.7% of the subjects complied with the request.
3. 25% agreed on the small request.

3.3 Repetition

The Premise The *repetition* principle has been developed and tested by Petty and Cacioppo [PC79]. Simply put, the technique calls for the reformulation of arguments and presenting them multiple times.

This repetition is not endless. Although it has been shown that repeating an argument two or three times, under different formulations, may be beneficial, this effect tends to decrease as the number of repetition increases.

Studies Petty and Cacioppo [PC79] conducted two studies in which participants heard the same argument (in different formulations) zero (for control), one, three or five times in succession. They were then asked to rate their agreement with the goal argument and list the arguments they could recall.

The results show that participants' agreement increases for the first three conditions and decreases when the argument is presented five times. The exact opposite pattern appears for the number of counter-arguments generated by the participants.

3.4 Anchoring

Principle The *anchoring* technique has been described by Tversky and Kahneman [TK75]. It refers to the tendency of people to generate judgments and estimations based on an initial reference point, an *anchor*. It is thus quite simple to manipulate this anchor by providing it in the

argument itself, for instance. In such case, subsequent judgement made by the persuadee are expected to be biased toward the anchor initially provided. While this technique is more commonly used in numerical settings (e.g., providing an anchor for a value of a product), it can also be applied to any type of arguments that could be ranked. In a sense, anchoring can be seen as a generalization of both the *foot in the door* and the *door in the face* principles. Indeed, in the former case, the persuader uses a smaller request first while in the latter she uses a bigger request first.

Studies Tversky and Kahneman [TK75] provided several examples of the use of the anchoring technique. In one of the studies reported, experimenters asked participants to generate a series of estimations (in percentages), such as the percentage of African countries in the United Nations (UN). Before the estimation was provided, a random number between 0 and 100 was presented to participants by spinning a wheel in the participants presence. The results show that different initial number presented on the wheel led participants to generate different estimations: the group that received the number 10 estimated that 25% of African countries were in the UN (on average), whereas the group that received the number 65 estimated that 45% of African countries were in the UN (on average).

In another study reported, experimenters asked two groups of students to estimate, within five seconds, the product of a numerical expression:

1. $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$, or,
2. $1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$, *i.e.*, the exact same sequence but in reverse order.

The median estimate for the first group (with the descending sequence) was 2250, whereas it was 512 for the second group (with the ascending sequence). This result is can be simply explained by the fact that subjects based their estimation on the results of the first operations they were able to make, which are obviously higher in the descending sequence.

4 From Psychosocial Persuasion Principles to Argumentation Frameworks

4.1 Foot in the Door

Argumentation Framework: Figure 2 depicts the application of the *foot in the door* technique to a very simple graph. Bold font arguments are the arguments from the persuader and the number represents the step in which each argument is played. **Zero means that the argument is not played at all.** In the original graph depicted in Figure 2a, the persuader wants to have argument **a** accepted at the end of the debate (it is the goal request). However, it is attacked by argument **b** from the persuadee and is thus not accepted according to classical semantics. Using the foot in the door principle, the persuader starts by playing a small argument **c** in order to have it accepted by the persuadee (argument




Fig. 2: Foot in the door

d signals acceptance). She then plays argument **a** as in the original graph. However, this time, the argument **a** is supported (indirectly) by argument **d** from the persuadee herself. This support is the attack on the original attacking argument (argument **b**), potentially preventing the persuadee from playing it. At the end, argument **a** is accepted.

Example 2. The arguments of Figure 2 can be instantiated as follows:

a Can we come to your house to ask you questions?

b No.

c Can you answer some questions over the phone?

d Ok.

The strategy of the persuader is to avoid the “No” by triggering the “Ok” to attack it.

The WASA associated to Figure 2a is:

$$\mathbb{A} = \left\langle \begin{pmatrix} a \\ b \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} w_a \\ w_b \end{pmatrix} \right\rangle$$

Analysis With $d = 2$, the propagation matrix after convergence is:

$$m_1 = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{1}{2} \end{pmatrix}$$

When calculating the degree of acceptability for both a and b , we have:

$$\begin{pmatrix} \frac{w_a - w_b}{2} \\ \frac{w_b}{2} \end{pmatrix}$$

This means that the initial plausibility for argument a needs to be at least twice as high as the one of argument b if we consider $d = 2$ as the real dampening factor. However, as we can see in the user studies, only 22.2% of the control group agreed with the request. Therefore, the initial strength for argument a is a fifth of the one of argument b .

When applying the same procedure on the WASA for Figure 2b:

$$\mathbb{A}' = \left\langle \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} \begin{pmatrix} 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} w_a \\ w_b \\ w_c \\ w_d \end{pmatrix} \right\rangle$$




Fig. 3: Door in the face.

The propagation matrix after convergence, again with $d = 2$ is:

$$\begin{pmatrix} 0 & -\frac{1}{2} & \frac{1}{4} & \frac{1}{8} \\ 0 & 0 & -\frac{1}{2} & -\frac{1}{4} \\ 0 & 0 & 0 & \frac{1}{2} \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

When multiplying, we can find that the degree of a is greater than the degree of b if the initial degree of d is at least one half bigger than the initial degree of b . In the experiments, as the disagreement with a is of 47.3% (and thus a is accepted), it would mean that about 24% at least where agreeing with answering the questions but not with the goal request.

4.2 Door in the Face

Argumentation framework In this case, the objective is to trigger a “No” response on an extreme request for the persuadee to feel somewhat obliged to not say “No” to the goal request.

Example 3. Let us instantiate the arguments of Fig. 3, in a context of the zoo trip experiment, as follows:

- a** Look after juvenile delinquents for a two-hour trip.
- b** No.
- c** Look after them for two years.
- d** No.

The WASA associated with Figure 3b is:

$$\mathbb{A}' = \left\langle \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} \begin{pmatrix} 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} w_a \\ w_b \\ w_c \\ w_d \end{pmatrix} \right\rangle$$

Interestingly, the nature of the relation between arguments c and d does not change the analysis. Indeed, as for the foot in the door technique, argument a is accepted (its plausibility is bigger than the one of its attackers) if the initial degree of d is at least one half bigger than the one of b .

4.3 Repetition




Fig. 4: Repetition

Figure 4 depicts how the repetition technique can be represented. Argument a is the goal request while arguments b to f are the **same** argument (in different formulations) repeated a certain number of times. As stated before, this differs from traditional abstract argumentation where arguments are represented in the graph irrespectively of the way they are used. In this case, by representing the timestep at which the argument has been played allows us to represent that an argument has been played several times in the dialogue. Argument g is a fictitious argument reinforced each time a repetition is made.

The WASA associated is as follows:

$$\mathbb{A} = \left\langle \begin{pmatrix} a \\ b \\ c \\ d \\ e \\ f \\ g \end{pmatrix} \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 \end{pmatrix} \begin{pmatrix} w_a \\ w_b \\ w_c \\ w_d \\ w_e \\ w_f \\ w_g \end{pmatrix} \right\rangle$$

Interestingly, with a dampening factor $d = 2$ we can see that the acceptability degree of a is bigger than the one of e as long as the sum of initial plausibilities for the repeating arguments is bigger than twice the one of e . However, we want to achieve a decrease in the degree of a

if we play more than three repeating arguments. A straightforward way to have this behaviour is to give a proper negative initial value to all repeating arguments after the third.

4.4 Anchoring

As a generalization of the foot in the door and the door in the face principle, the argument graph for the anchoring principle is also a generalization. However, the initial plausibilities in the WASA are no tied to the actual arguments this time but rather to their position in the ranking and the objective in the persuasion problem. For instance, if the objective is to sell a car at the highest price possible, the first argument should be a price above the actual price and then, in a second time, the actual price. On the other, if the objective is to buy the very same car, it is better to give a very low price first and then converge towards the price we were willing to pay from the beginning. Therefore, in the former case, the initial value should be low for the extremely high price, increasing the more it closes the gap with the price, and then decreasing again as it goes further down, past the price. It is the opposite in the latter case.

5 Discussion

In this paper we have made a modest step towards bridging the gap between abstract argumentation and psychological evidence for persuasion. This can be viewed as part of a larger effort to investigate what drives human decision-making in the argumentative context [RK18].

We have explained how to design argument graphs modelling four different psychological techniques, commonly used by people, and we have shown how they can be used to theoretically explain the observed results in human studies.

In future works we plan to include additional psychological techniques and a deeper analysis of the WASA and other abstract argumentation frameworks. A comparison with the traditional semantics in bipolar argumentation frameworks is another interesting direction. Finally, performing user studies based on the new representation is crucial to validate this new hybrid formalization.

References

- [ACLSL08] Leila Amgoud, Claudette Cayrol, Marie-Christine Lagasque-Schiex, and Pierre Livet. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems*, 23(10):1062–1093, 2008.
- [BCH17] Elizabeth Black, Amanda J Coles, and Christopher Hampson. Planning for persuasion. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 933–942. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

- [BPW14] Gerhard Brewka, Sylwia Polberg, and Stefan Woltran. Generalizations of dung frameworks and their role in formal argumentation. *IEEE Intelligent Systems*, 29(1):30–38, 2014.
- [BRT⁺15] Pietro Baroni, Marco Romano, Francesca Toni, Marco Auricchio, and Giorgio Bertanza. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation*, 6(1):24–49, 2015.
- [CTO14] Federico Cerutti, Nava Tintarev, and Nir Oren. Formal arguments, preferences, and natural language interfaces to humans: an empirical evaluation. In *Proceedings of the Twenty-first European Conference on Artificial Intelligence*, pages 207–212. IOS Press, 2014.
- [CVL⁺75] Robert B Cialdini, Joyce E Vincent, Stephen K Lewis, Jose Catalan, Diane Wheeler, and Betty Lee Darby. Reciprocal concessions procedure for inducing compliance: The door-in-the-face technique. *Journal of personality and Social Psychology*, 31(2):206, 1975.
- [DHM⁺11] Paul E Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence*, 175(2):457–486, 2011.
- [Dun95] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.
- [FF66] Jonathan L Freedman and Scott C Fraser. Compliance without pressure: the foot-in-the-door technique. *Journal of personality and social psychology*, 4(2):195, 1966.
- [Fog98] Brian J Fogg. Persuasive computers: perspectives and research directions. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 225–232. ACM Press/Addison-Wesley Publishing Co., 1998.
- [HBM⁺15] Emmanuel Hadoux, Aurélie Beynier, Nicolas Maudet, Paul Weng, and Anthony Hunter. Optimization of probabilistic argumentation with markov decision models. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [HH17] Emmanuel Hadoux and Anthony Hunter. Strategic sequences of arguments for persuasion using decision trees. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, 2017.
- [MN16] Till Mossakowski and Fabian Neuhaus. Bipolar weighted argumentation graphs. *arXiv preprint arXiv:1611.08572*, 2016.
- [PC79] Richard E Petty and John T Cacioppo. Effects of message repetition and position on cognitive response, recall, and persuasion. *Journal of personality and Social Psychology*, 37(1):97–109, 1979.
- [RK14] Ariel Rosenfeld and Sarit Kraus. Argumentation theory in the field: An empirical study of fundamental notions. In *Proceedings of the Workshop on Frontiers and Connections between Argumentation Theory and Natural Language Processing, Forlì-Cesena, Italy, July 21-25, 2014.*, 2014.

- [RK16a] Ariel Rosenfeld and Sarit Kraus. 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, December 2016.
- [RK16b] Ariel Rosenfeld and Sarit Kraus. Strategical argumentative agent for human persuasion. In *22nd European Conference on Artificial Intelligence (ECAI)*, volume 285, page 320. IOS Press, 2016.
- [RK18] Ariel Rosenfeld and Sarit Kraus. *Predicting Human Decision Making: From Prediction to Intelligent Agent Design*. Morgan & Claypool, 2018.
- [RMB⁺10] Iyad Rahwan, Mohammed I Madakkatel, Jean-François Bonnefon, Ruqiyabi N Awan, and Sherief Abdallah. Behavioral experiments for assessing the abstract argumentation semantics of reinstatement. *Cognitive Science*, 34(8):1483–1502, 2010.
- [TK75] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*, pages 141–162. Springer, 1975.