

Predicting persuasive effectiveness for multimodal behavior adaptation using bipolar weighted argument graphs

Klaus Weber, Kathrin Janowski, Niklas Rach, Katharina Weitz, Wolfgang Minker, Stefan Ultes, Elisabeth André

Angaben zur Veröffentlichung / Publication details:

Weber, Klaus, Kathrin Janowski, Niklas Rach, Katharina Weitz, Wolfgang Minker, Stefan Ultes, and Elisabeth André. 2020. "Predicting persuasive effectiveness for multimodal behavior adaptation using bipolar weighted argument graphs." In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '20)*, 1476–84. New York, NY: ACM.
<https://dl.acm.org/action/showFmPdf?doi=10.5555%2F3398761>.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



Predicting Persuasive Effectiveness for Multimodal Behavior Adaptation using Bipolar Weighted Argument Graphs

Klaus Weber
Human-Centered
Multimedia
Augsburg University
Augsburg, Germany
weber@hcm-lab.de

Kathrin Janowski
Human-Centered
Multimedia
Augsburg University
Augsburg, Germany
janowski@hcm-lab.de

Niklas Rach
Institute of Communication
Engineering
Ulm University
Ulm, Germany
niklas.rach@uni-ulm.de

Katharina Weitz
Human-Centered
Multimedia
Augsburg University
Augsburg, Germany
weitz@hcm-lab.de

Wolfgang Minker
Institute of Communication
Engineering
Ulm University
Ulm, Germany
wolfgang.minker@
uni-ulm.de

Stefan Ultes
Institute of Communication
Engineering
Ulm University
Ulm, Germany

Elisabeth André
Human-Centered
Multimedia
Augsburg University
Augsburg, Germany
andre@hcm-lab.de

ABSTRACT

Research states that persuasion is subjective. Moreover, people use behavioral cues all the time, very often even without noticing and are often not aware of being persuaded by non-rational cues. In order to draw attention to these effects, we want to enable virtual agents to adapt their behavior during interaction to the listener in order to increase their perceived power of persuasion.

In this paper, we introduce a novel multi-modal persuasive AI system that presents arguments from an underlying logical argument structure to a user by means of a virtual agent and synthetic speech. In doing so, the agent is able to adapt its multimodal behavior to the user, based on his or her explicit feedback. To this end, the feedback is used to predict the current user's stance by considering the underlying argument structure using bi-polar weighted argument graphs to later optimize the adaptation of the multimodal presentation by means of Reinforcement Learning.

We report on results of a user study with 48 participants showing the validity and practical potential of the proposed prediction model and conclude by providing limitations and implications in detail.

KEYWORDS

Adaptive Persuasive Agents, Behavior Policy Learning, Persuasion, Human-Agent-Interaction

ACM Reference Format:

Klaus Weber, Kathrin Janowski, Niklas Rach, Katharina Weitz, Wolfgang Minker, Stefan Ultes, and Elisabeth André. 2020. Predicting Persuasive Effectiveness for Multimodal Behavior Adaptation using Bipolar Weighted Argument Graphs. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 9 pages.

1 INTRODUCTION

The communication of opinions along with different pro- and counter-arguments is an important factor in the process of opinion building. Especially in political debates or public speeches, we often see that people without previous knowledge or opinion can be persuaded very easily by arguments even though they are not valid from a rational point of view. Presenting content-wise identical arguments in different ways, such as changing body language, gazing behavior, emotions or even the identity of the speaker can have a different effect on the audience's opinion towards this argument and overall stance regarding the topic.

The effect of an argument and whether or not it is perceived as persuasive does not only depend on the content itself but also the presentation of it, the credibility of the speaker and the own beliefs of the audience. Besides, in accordance with persuasive theory (see Sec. 2.1), it is not only important, *what* is said (semantic content), but also *how* something is said. This may include *verbal*, but also *non-verbal* techniques, such as gaze, gestures, use of emotions as well as rhetorical quality aspects, etc.

Wachsmuth et al. [36] state that argumentation quality aspects include *logical*, *rhetorical* and *dialectical* quality aspects. They further indicated that “*rhetorical quality is reflected by the persuasive effectiveness*”, which means persuading the *persuadee* successfully.

Studies with humans and robots have already shown that non-verbal cues affect the persuasive effectiveness [4, 7, 13, 31]. While most persuasive dialog systems refer to logical and verbal aspects (see Sec. 2.2), to our best knowledge, non-verbal aspects influencing effectiveness are fairly neglected in such systems.

However, Fogg et al. [12] claimed that, in order to successfully persuade individuals, three factors must be satisfied: 1) the right content, 2) at the right time as well as 3) the right strategy. In most systems, the first two aspects are mostly fixed [18], thus, researchers suggested that the adaptation of the way should be the main focus at the individual level [19]. In addition to that, O’Keefe and Jackson [22] pointed out that effectiveness is subjective. This

is in line with Kaptain et al. [19] who found that “*effectiveness of influence strategies varies from one person to another*”.

To build such a continuously adaptive system, we encountered the following key challenges: (i) Defining a model to correctly predict the user’s current stance considering the underlying argument structure and the feedback, (ii) enabling an agent to learn a behavior strategy that allows it to be perceived differently depending on its applied behavior and, (iii) allowing the agent for adapting to the audience in real-time based on explicit user feedback.

Within this work, we investigate a fine-grained model for predicting the user’s stance based on bi-polar weighted argument graphs and explicit user’s feedback and discuss a Reinforcement-Learning-based approach to utilize this information for real-time adaptation of the multi-modal behavior. Both aspects are evaluated in separate setups in order to verify their feasibility for the task at hand. Consequently, we include a user study with 48 participants to verify the predictive capability of the stance model as well as simulation results for the adaptation. To illustrate the adaptive feasibility of our prototype, the behavior of our adaptable agent is expressed through different emotional expressions (Section 2.1.2). Section 2 outlines research covering the theory of persuasion and non-verbal signals along with persuasion and argumentative persuasive systems in general. Section 3 gives an overview over the utilized argument structure while Section 4 describes our prototype and the proposed model in detail. In Section 5 and 6, we present the results of simulated users showing the adaptive feasibility, and a user study with 48 participants. Finally, we provide detailed lessons learned and insights into our future work of this research in Section 7.

2 RELATED WORK

Related research can be separated into two streams: (1) research focusing on persuasion and (2) research on persuasive systems:

2.1 Research Referring to Persuasion

2.1.1 Theory of Persuasion. The theory of persuasion goes back to Aristotle, who identified three means of persuasion: λόγος (logos), ἦθος (ethos) and πάθος (pathos) [21]. The term logos includes “[1.] *the process of identifying the issues at the heart of debate*; [2.] *the range of diverse arguments in the discourse*; [3.] *the structure of thoughts these arguments compose and* [4.] *the sequencing, coherence and logical value of these arguments*”. [9, p.18].

According to Wachsmuth et al. [36], several aspects need to be considered in order to make arguments persuasive in general, that are *logical quality* (including cogency, fallaciousness, and the argument’s strength), *rethorical quality* (including effectiveness) and *dialectical quality* (including convincingness, reasonableness, and global sufficiency). Thus, *logos* and *logical quality aspects* refer to *what* is said (semantic content), whereas *rethorical quality* refers to *how* something is said (the ways to express). The latter is covered by ἦθος (ethos) as well as πάθος (pathos), where ethos includes *personality* and *stance* [9], and pathos describes the emotional engagement between the *persuader* and the *persuadee*, which means appealing to the *persuadees’* (listener’s) emotions.

Psychological persuasion models developed by Petty and Cacioppo [23] (Elaboration Likelihood Model – ELM) and Chaiken et al. [6] (Heuristic-Systematic Model – HSM) describe the influence

of information processing on the result of persuasive messages. A persuasive message can be processed via two different cognitive routes, namely central and peripheral processing. Central processing focuses on the content of the message communicated by the *persuader*/agent, while peripheral processing focuses on the expression of the agent. However, people do not process information in isolation via the central or peripheral route [6]. Rather, peripheral processing is always carried out, in addition to which, if an elaboration threshold is reached, central processing also takes place. In this situation, the two processing paths are used with different intensities depending on the audience’s “*need for cognition*” [23]).

In summary, persuasion is a complex process consisting of more than just logical aspects. Thus, the development of persuasive systems is a difficult task considering all these aspects. The complexity of these along with the subjectivity of effectiveness motivates the need for adaptation approaches to enhance the persuasive effectiveness of virtual agents autonomously during interaction.

2.1.2 Non-verbal Signals and Persuasion. Non-verbal signals and their impact on persuasion have been extensively explored by other researchers. It has been shown that robots using vocal cues and behavior language, such as gaze and gestures, tend to be more persuasive than robots not using them. Chidambaram et al. [7] compared the effect of non-verbal cues (gesture, gaze, etc.) to vocal cues showing a higher persuasive effect for non-verbal cues. Ham et al. [13] showed that communication through gaze led to a higher persuasive effect than gestures did. Siegel et al. [31] investigated the effect of gender aspects, while Andrist et al. [4] compared the persuasive effect of practical knowledge and rhetorical ability. DeSteno et al. [10] showed that persuasive messages were more successful when framed with emotional overtones appropriate to the emotional state of the recipient. They say that “*emotions can alter the persuasive impact of messages as a function of the emotional framing of persuasive appeals*”. Wang et al. [37] demonstrated the effect of emotions depending on the level of power of expresser and recipient, whereas Van Kleef et al. [35] showed that people use the source’s emotions as information channel when forming their own attitudes, i.e. different emotions led to different attitudes. This finding is in line with the *EASI* theory [34], which states that individuals are influenced by the emotional state of others. There is a lot of evidence from the literature that the perceived persuasiveness of arguments depends largely on emotions. We, therefore, felt encouraged to focus on emotions as a major component of our adaptation approach. However, the herein discussed models do not depend on a specific multi-modal behavior.

2.2 Persuasive Systems

Hunter states that “*An automated persuasion is a system that can engage in a dialogue with a user (the persuadee) in order to persuade to do (or not do) some action or to believe (or not to believe) something*” [16]. A lot of research on persuasive systems focuses on the logical representation of arguments with the initiative solely belonging to the system. Systems modeling the argumentative interaction as a dialog game were introduced in [5, 40] and rely on heuristics or predefined strategies for determining the next actions (i.e., arguments). An argumentative dialog system for the exchange of arguments with a speech interface was introduced in [15].

Different setups utilizing Reinforcement Learning to optimize the system logical strategy were introduced: Rosenfeld and Kraus [30] described a persuasive agent that learns the optimal response from a recorded corpus of arguments and employs a Bipolar Weighted Argument Framework which only allows the exchange of arguments. Alahmari et al. [2] investigated Reinforcement Learning in the context of the dialog game. Multi-agent Reinforcement Learning in the context of argumentation was addressed by Rach et al. [25]. The authors included the derived strategies into a multimodal dialogue system [26] with no adaptation of the behavior, though.

“Project Debater”, an advanced persuasive system by IBM¹, can debate humans on complex topics. It is based on three pioneering capabilities: data-driven speech writing, listing comprehension as well as modeling of human dilemmas. Their system, however, only provides speech output and completely neglects multi-modal output, such as body language, as well as adaptation.

All of these listed works focus solely on the logical part of the strategy, hence only addressing the question of *what* to say whereas, in contrast, the presented work also considers the question of *how* to present an argument depending on the user’s preferences.

Another form of persuasive systems concerns recommendation systems and personalized presentations of recommendations. For instance, Hammer et al. [14] enabled systems to present personalized recommendations in order to increase the perceived persuasiveness of recommendations. Yet, no continuous adaptation was applied.

Ritschel et al. [29] developed a robotic health companion giving health-related recommendations. In doing so, the robot dynamically adapts the politeness of its linguistic style based on explicit user feedback. However, the recommendations only consisted of single term recommendations and were not, contrary to our approach, based on argumentation structures.

The only and (to our best knowledge) first persuasive argumentation system we are aware of using a continuous adaptation approach based on users’ preferences was introduced by Kang et al. [18]. Their approach is based on the aforementioned ELM. They enabled an agent to determine which route (central or peripheral) of thinking is most likely used by the user based on the current mental user state consisting of their *ability* as well as *motivation*. Overall, their system consists of eleven different verbal strategies.

While their approach is able to learn a persuasive strategy based on the most likely used route of thinking, they did not take into account, in contrast to our approach, the underlying argument structure, which is important to realize the full persuasive potential of such systems. Further, they only used strategies in general, but did not explicitly consider the influence of multi-modal strategies.

3 ARGUMENT STRUCTURE AND NLG

The arguments that are available to the system throughout the interaction are encoded in an argument structure based on the argument annotation scheme introduced in [32]. It includes three types of argument components (*Major Claim*, *Claim*, and *Premise*) and two different directed relations (*support* and *attack*) between them. Relations are allowed from *Claims* to the *Major Claim*, *Premises* to the *Major Claim*, *Premises* to *Claims* and *Premises* to *Premises*. If a component φ_i has a relation towards a component φ_j , we say that φ_j is

the target (of φ_i) and each component (apart from the *Major Claim* φ_0) has exactly one target. Hence, the arguments $\mathbf{arg}_i \in \mathbf{Args}$ that can be generated from such a structure have the form $\mathbf{arg}_i = (\varphi_i \text{ so } \varphi_j) (\hat{=} \text{support})$ or $\mathbf{arg}_i = (\varphi_i \text{ so } \neg\varphi_j) (\hat{=} \text{attack})$. Since each relation is unique and the difference between the three types of components is characterized solely by the allowed relations, each resulting structure can be represented as acyclic directed graph with argument components as nodes and relations as edges (see Figure 1). For the sake of simplicity, the *Major Claim* is defined as $\mathbf{arg}_0 = \varphi_0$ throughout this work.




Figure 1: Argument structure consisting of components φ_i and φ_j as well as logical statements between them.

In addition, each argument $\mathbf{arg} \in \mathbf{Args}$ refers to one of two existing *stances* $\in \{+, -\}$ of the topic. For the sake of simplicity, the stance of $\mathbf{arg}_0$ is considered as $+$ throughout this work. The other arguments’ stances are computed with respect to the arguments’ relations, i.e., the stance of supporting arguments is always the same as its target’s stance, while the stance of attacking arguments is the opposite of the target’s one.

Throughout this work, we aim for *non-opinion-based* topics, i.e., topics for which we can assume a minimal bias of the user, as our goal is to investigate the effect of the agent’s non-verbal behavior. To this end, we utilize hotel reviews from the annotated *SemEval-2015 Task 12* Test Datasets [24] and infer the argument structure for each hotel from the labels with a procedure adapted from the argument mining approach described in [8]. This choice is due to the fact that the corpus provides high-quality annotations on a large data set that includes all the required information:

- An aspect category E#A consisting of an entity E (e.g. Hotel, Service, Location) and an attribute A (e.g. Price, Quality).
- A polarity (positive, negative, neutral) of the category.
- An opinion target expression within the annotated sentence that explicitly refers to the entity (if present).

The annotation is illustrated by the following example that includes the original sentence as well as the annotated labels: *Vending machines were out of everything except in the lobby.* (target: *Vending machines*, category: FACILITIES#GENERAL, polarity: *negative*)

The argument structures can be derived from this annotation as follows. For each hotel, we generate a template consisting of the *Major Claim* “This hotel is worth a visit” and one *Claim* for each entity E. The polarity of each *Claim* is determined by comparing the number of positive and negative sentences annotated with the respective entity. All sentences that have consistent polarity annotations for at least one entity, are not marked as “OutOfScope”, and have an opinion target expression are then included into the structure. To this end, we assume that sentences within one review and with the same entity label build on each other, which

¹<https://www.research.ibm.com/artificial-intelligence/project-debater/>

means that the first sentence targets the respective *Claim* and the following build a chain of arguments. For the remaining argument components, we assume a direct relation to the respective *Claim* unless the component in question has the same target expression as a predecessor, in which case it is directly related to the same. Finally, sentences that include multiple labels related to different *Claims* are included in all related sub-structures and (if possible) manually separated into different arguments. In all cases, the type of relation is determined by the polarity.

The annotated sentences serve as a template for the Natural Language Generation (NLG) of the system. Grammatically incomplete arguments were revised manually to form a stand-alone sentence and repeated arguments were merged into one. In addition, connecting phrases between the arguments were included in a separate template in order to ensure a fluent interaction. These additional statements include a notification about the stance of the presented argument (if the current argument attacks its target), a notification if the presented argument refers not to the immediate predecessor in the conversation, as well as an introduction and closing statement. Apart from the introduction and closing statement, multiple formulations for each case were included in the template from which the system selects randomly for each utterance. The following exemplary utterance includes a topic switch (ts1), the referenced argument (a1), a notification about the stance (s1) and the new argument (a2): "The next argument is related to something I mentioned earlier. I said (ts1): All in all, it is a nice and affordable spot for sightseeing in the area (a1). I also found an opinion that disagrees with this aspect. The respective author wrote (s1): I think all in all the price was way too high for such a poor accommodation (a2)."

4 ADAPTATION MODEL AND PROTOTYPE

In the following, we sketch the design of our prototype and, afterwards, present the adaptation model in detail.

4.1 System overview




Figure 2: Prototype of our web interface consisting of an agent presenting her arguments to a user along with emotional behavior. The user gives feedback (*convincing*, *neutral*, *not convincing*) about the agent’s persuasive effectiveness after each argument, which is used to train the strategy.

To allow users for interacting with our system, we extended the web-interface of [26] as visualized in Figure 2. The agent’s task is to talk about a topic (*Major Claim*, see Section 3) by giving arguments that are either *for* or *against* the topic. To do so, the system first assigns a *stance* to the agent, which is either *pro* (+) or *con* (-). Then, the agent presents both pro- and counter-arguments about the topic in different ways to convince the user to either change their belief (if the user is against the agent’s stance) or to increase their beliefs towards the corresponding stance. To influence the user, the agent can underlie its argument with appropriate emotions.

At every interaction step, a random argument $arg \in Args''$ is selected from all possible next arguments $Args'' \subset Args$. The subset $Args''$ denotes all arguments that have a direct link to any argument already given before, denoted as $Args' \subset Args$. Then, an appropriate emotion is selected depending on the agent’s policy π and presented to the user in combination with the selected *arg* (Figure 3).




Figure 3: Conceptual model: The agent selects an argument $arg \in Args'' \subset Args$ connected to any previously given argument $Args' \subset Args$ as well as an emotion according to the learned policy π and presents it to the user.

After each argument, the user provides the agent with explicit feedback by using the feedback buttons (*convincing*, *neutral*, *not convincing*) as illustrated in Figure 2. This feedback is subsequently used to determine the agent’s persuasive effectiveness (Definition 4.2) and to learn the policy π that works best for the current user.

4.2 Multimodal Output



Figure 4: Example emotions: Angry, happy, and sad (FLTR).

To present the arguments via multimodal output to the user, we employ the Charamel 3D character rendering engine². In this work, we use the Gloria avatar capable of performing social-cue based interactions with the user as illustrated in Figure 2. The avatar can perform lip-sync speech output using the Nuance TTS along with the Amazon Polly voices. Further, the avatar comes with more than 50 motion-captured gestures as well as a set of facial expressions

²<http://www.charamel.com>

(14), including the basic emotions defined by Ekman [11]. Some examples of the avatar’s possible emotions are given in Figure 4.

4.3 Adaptation and Prediction Model

In the following section, we describe the RL method including the employed prediction model in detail. As stated by Ritschel et al. [28] and Weber et al. [38] an adaptation to human preferences should work during interaction. Thus, some simplifications are needed for our system in terms of state and action space.

To enable the agent to learn the user’s preferences (policy π) during interaction, we apply Reinforcement Learning (RL) based on linear function approximation along with a basis transformation using the Fourier Basis as proposed by Konidaris et al. [20]. This approach has three advantages: (1) Reinforcement Learning allows for learning a sophisticated policy π based on trial and error (during interaction). (2) A linear function approximator bears the advantage that multiple similar states can be learned at the same time while allowing for quick adaptation compared to non-linear methods. (3) Using the Fourier Basis allows for learning explicit linear and non-linear dependencies between state-input parameters.

4.3.1 State Space. Every RL state $s \in \mathcal{S}$ is generated based on the current argument $arg \in \mathcal{Args}$. As aforementioned in Section 3, each argument has a *stance* $\in \{+, -\}$ it refers to as well as a *relation* $\in \{\text{attack}, \text{support}\}$. To enable the agent to further optimize its behavior with respect to the conveyed sentiment of the argument, we use the sentiment analysis tool *vaderSentiment*³ by Hutto et al. [17] to compute information about the negativity, neutrality and positivity and the respective compound score of an argument. A state $s_{arg} \in \mathcal{S}$ for an argument $arg \in \mathcal{Args}$ is defined as:

Definition 4.1 (State). Let $\text{stan} : \mathcal{Args} \rightarrow \{+, -\}$ be the stance and $\text{rel} : \mathcal{Args} \rightarrow \{\text{attack}, \text{support}\}$ be the relation of an arbitrary argument $arg \in \mathcal{Args}$. Further let $\text{score} : \mathcal{Args} \rightarrow [-1, 1]$ be the normalized compound score of the sentiment analysis of an argument. Then, the state s_{arg} is defined as a 3-tuple

$$s_{arg} := (\text{stan}(arg), \text{rel}(arg), \text{score}(arg)) \quad (1)$$

4.3.2 Action Space. As described above, we make use of the agent’s provided emotions. Thus, the discrete action space $\mathcal{A}$ consists of different emotions (facial expression and gestures), such as *happy*, *sad*, *angry*, *disappointed*, etc., that can be displayed by the agent with different discrete intensities $\in \{\text{low}, \text{medium}, \text{high}\}$.

4.3.3 Prediction Model and Reward Function. Defining an appropriate reward signal can be challenging. As aforementioned, the user gives feedback $f(arg)$ for every $arg \in \mathcal{Args}$ (i.e., state s_{arg}) with $f : \mathcal{Args} \rightarrow \{1.0, 0.5, 0.0\}$, which translates to:

- *Convincing*, i.e., positive feedback ($f = 1.0$)
- *Neutral* ($f = 0.5$)
- *Not convincing*, i.e., negative feedback ($f = 0.0$)

Considering the argumentation structure (acyclic directed graph with two types of relations, see Section 3), one cannot directly take these feedback signals as reward since we have to consider the goal of the agent, that is, as mentioned before, convincing the user of the assigned *stance* ($+$, $-$). That means, if the user finds an argument

convincing, we have to check the *stance* it belongs to and invert the feedback signal $x = f(arg)$ using $\text{inv}_{arg}(x)$ (see Equation 2) respectively, if the argument’s stance does not match the agent’s stance, i.e., $\text{stan}(arg) \neq \text{assigned stance}$.

$$\text{inv}_{arg}(x) = \begin{cases} 1.0 - x & \text{if } \text{stan}(arg) \neq \text{assigned stance} \\ x & \text{else} \end{cases} \quad (2)$$

Even though, in theory, using $\text{inv}_{arg}(x)$ as a reward would generally allow the agent for adaptation, this approach does not allow the agent to observe to what extent the goal of convincing the user of the *assigned stance* could already have been achieved compared to the *non-assigned stance*. Further, not all arguments have the same impact on the overall effectiveness of the *Major Claim* (topic) because argument components are linked to each other and, thereby, influence the effectiveness of others. Therefore, it seems fairly unnatural that an argument component that has many nodes between itself and the Major Claim has the same effect as argument components that are directly linked to the Major Claim.

Thus, the opinion of the user does not only depend on how many arguments could be weakened/strengthened by other arguments but also on the relations between the arguments. Consequently, it would come in handy to have a reward signal giving a prediction of the user’s current stance by considering these argument relations and structure, which can be done using reward shaping [39].

Prediction Model. To define such a prediction model, we got inspired by the work of Aicher et al. [1]. They recently presented an interactive system to assist users in their opinion-forming process by allowing them to give information about what arguments they prefer over others and which they reject. They then computed the user’s preferences based on bipolar weighted argument graphs (BWAGS) by utilizing a linear Euler-based restricted semantics introduced by Amgoud et al. [3]. BWAGS are generally used for computing the strength of arguments (in an acyclic directed argument graph) considering their own weight ω and the strengths s of their child arguments [3] as visualized in Figure 5.




Figure 5: Sketch of a BWAG for an arbitrary argument $m \in \mathcal{Args}$ having two children k and $k + 1$. The strength s_m for argument m is computed by a function g considering its own weight ω_m and the strength of the child nodes s_k and s_{k+1} .

Yet for our purpose, the Euler-based restricted semantics is not suitable, as arguments with a weight equal to zero are considered invalid and, thus, their children have no effect on the overall stance (see [1], [3]). However, using this information to correctly predict the user’s stance and to compute a reward signal, we are still interested in arguments having a persuasive effect of zero as this

³<https://github.com/cjhutto/vaderSentiment>

affects the overall persuasive effect of the agent negatively (or positively, depending on the argument's stance). Hence, we propose the following definition of the persuasive *effectiveness* e_i for an argument $i \in \text{Args}$ for predicting the user's current stance, which avoids considering arguments as invalid having a weight $\omega = 0$.

Definition 4.2 (Effectiveness). Let $i \in \text{Args}$ be an arbitrary argument, and let $\mathbb{A}_i \subset \text{Args}$ be all child nodes of i , i.e., all arguments having i as direct target. Further, we denote $\mathbb{A}_i^+ \subseteq \mathbb{A}_i$ as the subset of all supporting and $\mathbb{A}_i^- \subseteq \mathbb{A}_i$ as the subset of all attacking arguments of argument i , respectively. We then compute the effectiveness level $e_i \in [0, 1]$ of argument $i \in \text{Args}$ recursively as:

$$e_i = \frac{\omega_i + \sum_{j \in \mathbb{A}_i} \iota^{-1}(e_j)}{1 + |\mathbb{A}_i|} \quad (3)$$

where $\iota^{-1} : [0, 1] \rightarrow [0, 1]$ defines the inverse function of e_j as

$$\iota^{-1}(e_j) = \begin{cases} e_j & \text{if } j \in \mathbb{A}_i^+ \\ 1 - e_j & \text{else} \end{cases} \quad (4)$$

and where

$$\omega_i = \begin{cases} f(i) & \text{argument } i \text{ used by agent} \\ 0.5 & \text{else} \end{cases} \quad (5)$$

Please note, that the effectiveness level e_j is computed with respect to the argument's stance, not with respect to the agent's assigned stance, i.e., we only have to invert the effectiveness levels of all arguments $j \in \mathbb{A}_i^-$, i.e., all attacking arguments following the general idea of bipolar weighted argument graphs.

As seen in equation 5, each weight in the argument graph is initialized with 0.5, which consequently yields $e_0 = 0.5$. The effectiveness level is re-computed from bottom to top recursively after every weight change. It is easy to verify that neutral user feedback ($f=0.5$) does not change the effectiveness level e_0 .

The computed effectiveness can be used to get a prediction of the user's current stance using Definition 4.3.

Definition 4.3 (Prediction). Let e_0 be the current computed effectiveness, then the user's stance is predicted as follows:

$$\text{stan}_{\text{user}} = \begin{cases} + & e_0 \geq 0.5 \\ - & \text{else} \end{cases} \quad (6)$$

It is worth mentioning that no interval is defined for an *unknown* stance. Instead, the effectiveness e_0 can be used to determine the confidence value of how sure the agent is about the prediction by looking at how close the effectiveness is to the criterion value 0.5.

We, finally, define the reward $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ based on the predicted stance as:

Definition 4.4 (Reward Function). Let $s_t \in \mathcal{S}$ be the current state and $a_t \in \mathcal{A}$ an action at RL time step t , and let $e_{0,t}$ be the current effectiveness after performing action a_t as well as $e_{0,t-1}$ the previous one, then the reward $\mathcal{R}_t(s_t, a_t)$ at time step t is defined as:

$$\mathcal{R}_t(s_t, a_t) := n \cdot (\text{inv}_0(e_{0,t}) - \text{inv}_0(e_{0,t-1})) \quad (7)$$

Initial tests have shown that $n \approx |\text{Args}|$ works best. Equation 2 is applied because for the main stance always yields $\text{stan}(0) = +$ and, thus, we have to compare the *assigned stance* with $\text{stan}(0)$ and invert the effectiveness level accordingly.

4.3.4 Algorithm. At every time step t , the agent selects one of the available actions $a \in \mathcal{A}$ according to the current state $s_t \in \mathcal{S}$ for argument $\text{arg} \in \text{Args}$ and its policy π (ϵ -greedy with $\epsilon = 0.05$) and uses the obtained feedback signal $f(\text{arg})$ to compute the current effectiveness level $e_{0,t}$ and reward signal $\mathcal{R}_t$.

Since we use linear function approximation, the agent has to learn a weight vector ω for every RL action $a \in \mathcal{A}$. The weight vector allows the agent to compute a value $Q(s_t, a_t, \omega)$ for every action $a_t \in \mathcal{A}$ and $s_t \in \mathcal{S}$ by calculating the dot product of the vector ω and vectorial representation of the current state $\phi(s_t)$, i.e.,

$$Q(s_t, a_t, \omega) := \phi(s_t) \circ \omega, \forall s_t \in \mathcal{S}, \forall a_t \in \mathcal{A} \quad (8)$$

The agent uses the reward $\mathcal{R}_t$ to update the weight vector ω_t until the policy π converged to the optimal one π^* (Sutton et al. [33]):

$$\Delta \omega_t = \alpha (\mathcal{R}_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}, \omega_t) - Q(s_t, a_t, \omega_t)) \phi(s_t) \quad (9)$$

5 INITIAL SIMULATION

In order to demonstrate the general adaptive feasibility of the Reinforcement Learning approach and to show that our system can handle different types of users, we tested our system in a simulation setup. We, thus, have run a simulation of 1000 simulated users with different behavior preferences as well as stances (either *for* or *against* the agent's stance) that work best to persuade them individually, e.g., users where the agent was the most effective when looking sad while presenting negative arguments that are against its own stance, and happy when presenting positive arguments that are in favor to the agent's stance and vice versa. Some simulated users preferred higher intensities of the emotions, while others did prefer lower ones. Further, we varied the agent's *assigned stance* ensuring that all possible combinations were tested.

Since deterministic user feedback is far from realistic [27], we have run the same simulation with different levels of noise (5%, 10% and 25%), where noise simulates random feedback from the user, which not necessarily matches the optimal policy π^* .




Figure 6: Simulation results including 95% confidence intervals showing the cumulative effectiveness level over time with respect to the assigned stance showing a continuous increase over time even for high noise.

Figure 6 depicts the results showing the effectiveness levels over time with different degrees of noise where the shaded areas depict a 95% confidence interval. The results reveal that the effectiveness levels increased over time, as the agent managed to gradually move the user towards its assigned stance by learning the most effective behavior strategy. Even though the performance decreases due to non-deterministic feedback when including noise, it shows that the agent is still able to cope with such non-deterministic feedback and find the strategy that increases the effectiveness level.

6 USER STUDY

To evaluate the validity and accuracy of the proposed prediction model, i.e., to verify that our model is able to correctly predict the user’s stance, we conducted a between-subject user study.

6.1 Participants, apparatus, and procedure

We recruited 48 participants (32 male, 16 female, 18-30 years old) from a university campus. All participants were students. At the start of the study, they were informed about the general procedure and asked to provide the agent with feedback using the interactive system of Figure 2. After the session, they were asked, whether or not they like to visit the hotel. To avoid bias effects beforehand, they were not told about the system’s overall goal to predict their decision but asked to provide feedback on whether or not they find an argument helpful. Figure 7 depicts the general study setup showing a participant interacting with the agent.




Figure 7: Evaluation setup with our interactive agent.

In each session, the agent presented 43 arguments for and against the hotel, which took about 10-15 minutes. The agent computed the effectiveness e_0 based on the given feedback $f(arg)$ for every presented argument $arg \in Args$. The agent’s assigned stance was counter-balanced, i.e., half of the participants were interacting with an agent who was in favor of visiting the hotel and vice versa. During the study, we collected the following data:

- (1) Directly given user feedback $f(arg)$, $\forall arg \in Args$.
- (2) Computed effectiveness e_0 using the given feedback.
- (3) Subjective decision if users like to visit the hotel (post-study).

6.2 Results

In the following, we first plot collected data to explore trends in data and afterward present statistical tests.




Figure 8: Collected data: Effectiveness e_0 , percentage of pos. (neg.) feedback f^+ (f^-) with respect to the assigned stance.

6.2.1 General trends. Figure 8 summarizes the results for all participants depicting the agent’s final effectiveness e_0 , the percentage of *user feedback* in favor of the agent’s assigned stance f^+ and the percentage of *user feedback* not in favor of the agent’s assigned stance f^- . Neutral feedback is not depicted as it does not affect the effectiveness (see 4.3.3). First, we notice two trends:

- (1) The higher (lower) the positive feedback, the higher (lower) the effectiveness e_0 .
- (2) The lower (higher) the negative feedback, the higher (lower) the effectiveness e_0 .

Thus, the positive feedback f^+ seems to positively correlate with the effectiveness and the negative feedback f^- seems to negatively correlate with the effectiveness e_0 . As stated, the general idea of the effectiveness e_0 is to get a prediction of the user’s current stance. So, we expect that a lot of positive feedback increases the effectiveness, while a lot of negative feedback decreases the effectiveness. The trends, therefore, are in line with our expectations.

6.2.2 Statistical Analysis. To verify the expected trends statistically, we computed the correlation between feedback and effectiveness showing a very strong and significant correlation (*positive correlation for positive feedback, negative correlation for negative feedback*).

	n	r	p
Pos. Feedback & effectiveness - Pearson correlation	48	0.92	<.001 ✓
Neg. Feedback & effectiveness - Pearson correlation	48	-0.83	<.001 ✓

Table 1: Correlation between feedback and effectiveness.

6.2.3 Prediction Accuracy. We then evaluated to what degree the predicted user’s stance $stan_{user}$ (see Definition 4.3) and the subjective user’s decision to visit the hotel match. We computed the agent’s confidence in the predictions based on how close they were on the criterion value ($e_0 = 0.5$). To this end, we utilize a modified sigmoid function to a) ensure that the extreme values of 0 and 1 correspond to a confidence of 100% and b) obtain a more fine-grained prediction for the most common interval $[0.3, 0.7]$. Consequently, a confidence $\geq 80\%$ means $e_0 \geq 0.64$ or $e_0 \leq 0.36$. The results in Figure 9 show that the objective system’s prediction is very accurate even for low confidence values proving the practical potential of the model.




Figure 9: Model accuracy of predicted user’s stance depending on different confidence values.

To verify the sensitivity and precision of the predictions, we computed the F1 score for both the positive and negative stance

depending on the prediction confidence as summarized in Table 2. The results show that the F1 score increases with higher confidences, thus, proving both the sensitivity and precision of the prediction.

Stance	Confidence			
	≥ 60%	≥ 80%	≥ 85%	≥ 90%
+	0.67	0.73	0.77	0.86
-	0.62	0.70	0.80	0.89

Table 2: F1 score for different prediction confidences.

7 DISCUSSION

In the beginning, we have argued that persuasion is subjective and, thus, an adaptation of behavior leads to an increased persuasive effectiveness of virtual agents. To this end, we have presented a model of a dynamically adaptive virtual agent capable of learning a behavior strategy during interaction with a human to increase its perceived persuasive effectiveness. We have further argued that it is important to consider the underlying argument structure to define a reward signal, which allows the agent to compute the degree of its persuasive effectiveness. Hence, we have presented a model that is able to predict the user’s current stance based on the underlying argument structure and bi-polar weighted argument graphs, which serves as the basis for the adaptation process. The adaptive feasibility could be proven in a carried out simulation beforehand. However, to justify the usage of the underlying prediction model for our adaptation process, we presented a thorough evaluation of the model. In a user study, conducted with 48 participants, we have found a very strong and significant correlation between feedback and effectiveness and have shown that our system is able to correctly predict the user’s stance proving the validity and practical potential of the model.

Argument structure as information source. In the beginning, we argued that different arguments can have different effects on the current effectiveness and, thus, the user’s stance based on their position in the argument structure. Despite the observed correlation between overall feedback and effectiveness at the end of the interaction, using bi-polar argument graphs as prediction model and reward function bears several advantages in comparison to models based solely on the statistics of the feedback: (i) The feedback for different arguments is weighted differently in the graph and, hence, differently affects the behavior strategy during the interaction, (ii) additional argument-specific or structure-specific information can be used in combination with the feedback to provide more fine-grained information for learning and (iii) behavior learning can be combined with fine-grained logical strategies [25, 30].

Predictive power of the model. It should be noted that the given user’s feedback during the study does not necessarily lead to their final decision for or against the hotel. Thus, it would not have been surprising if the predictions of the model were incorrect. However, the fact that we are able to predict the user’s current stance using the provided feedback in an agent-user interaction opens a lot of new possibilities in establishing successful persuasive systems. For instance, the predicted stance can be used to determine

when the user is likely convinced and the system can stop the persuasion process. Secondly, in persuasive debates between several agents/humans, the predicted stance can be used to determine the success of the whole debate during interaction and, thus, enables the agents to adopt strategies of the agent that is more successful.

8 LIMITATIONS AND FUTURE WORK

Even though we have proven the validity and high predictive sensitivity as well as precision, our approach still faces limitations that should not be neglected. The impact that the user’s feedback has on the effectiveness depends on the entire structure of arguments and, most importantly, on the number of arguments targeting any single argument. However, in the current version of the system, the feedback of arguments targeting the same parent argument arg_n (also known as siblings) is still equally weighted (unlike arguments targeting different arguments). It should be investigated whether this approach also works if there exists a too strong imbalance between arguments regarding their strength. A possible solution would be to enable users to provide additional information about how strong their feedback should influence the overall effectiveness compared to the sibling arguments or interval-scaled feedback could be used, similar to Aicher et al. [1]. In our future work, we aim to build on our results and explore the impact of adapting emotions expressed through multiple channels (facial emotions, gestures, and voice) in an extended user study that combines all discussed approaches.

9 CONCLUSION

We presented a novel multi-modal persuasive AI system with a virtual agent capable of presenting arguments from an underlying logical argument structure to a human user through synthetic speech and multimodal behavior. As research states that persuasive effectiveness is subjective, we enabled the agent to adapt its behavior to the human on the granularity of single arguments to increase its perceived persuasive effectiveness. To this end, we employ the user’s provided feedback to predict the user’s stance taking into account the underlying argument structure. In a detailed simulation, we showed the practical potential and the feasibility of our prototype. We have seen that the agent is able to adapt to a user during interaction, even for non-deterministic feedback. We conducted a user study to evaluate the validity and practicability of the underlying prediction model. The validation of the underlying model showed a significant correlation between feedback and computed effectiveness level and that, using this model, the system is able to correctly predict the user’s stance with high accuracy, even for low confidences making this model a powerful tool for persuasive systems. We have described lessons learned in detail including the predictive power of the model. We hope those insights will help fellow researchers in addressing future challenges in developing multimodal persuasive systems.

ACKNOWLEDGMENTS

This work has been funded by the Deutsche Forschungsgemeinschaft (DFG) within the project "How to Win Arguments - Empowering Virtual Agents to Improve their Persuasiveness", Grant Number 376696351, as part of the Priority Program "Robust Argumentation Machines (RATIO)" (SPP-1999).

REFERENCES

- [1] Annalena Aicher, Niklas Rach, Wolfgang Minker, and Stefan Ultes. 2019. Opinion Building based on the Argumentative Dialogue System BEA. In *Proceedings of the 10th International Workshop on Spoken Dialog Systems Technology (IWSDS 2019)*. Springer.
- [2] Sultan Alahmari, Tommy Yuan, and Daniel Kudenko. 2019. Reinforcement Learning for Dialogue Game Based Argumentation. In *Accepted of the 19th Workshop on Computational Models of Natural Argument (CMNA19)*.
- [3] Leila Amgoud and Jonathan Ben-Naim. 2018. Weighted Bipolar Argumentation Graphs: Axioms and Semantics. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden, 5194–5198. <https://doi.org/10.24963/ijcai.2018/720>
- [4] Sean Andrist, Erin Spannan, and Bilge Mutlu. 2013. Rhetorical robots: Making robots more effective speakers using linguistic cues of expertise. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, Tokyo, Japan, 341–348. <https://doi.org/10.1109/HRI.2013.6483608>
- [5] Trevor JM Bench-Capon. 1998. Specification and implementation of Toulmin dialogue game. In *Proceedings of JURIX*. Vol. 98. 5–20.
- [6] Shelly Chaiken, Akiva Liberman, and Alice H Eagly. 1989. *Heuristic and Systematic Information Processing within and beyond the Persuasion Context*. Guilford, 212–252.
- [7] Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. 2012. Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction - HRI '12*. ACM Press, Boston, Massachusetts, USA, 293–300. <https://doi.org/10.1145/2157689.2157798>
- [8] Oana Cocarascu and Francesca Toni. 2016. Detecting deceptive reviews using argumentation. In *Proceedings of the 1st International Workshop on AI for Privacy and Security*. 1–8.
- [9] Robert Cockcroft and Susan Cockcroft. 2003. *Persuading people: An introduction to rhetoric* (2nd ed.). Macmillan International Higher Education.
- [10] David DeSteno, Richard E Petty, Derek D Rucker, Duane T Wegener, and Julia Braverman. 2004. Discrete emotions and persuasion: the role of emotion-induced expectancies. *Journal of personality and social psychology* 86, 1 (2004), 43.
- [11] Paul Ekman. 1992. An argument for basic emotions. *Cognition & emotion* 6, 3-4 (1992), 169–200.
- [12] B.J. Fogg. 2007. *Mobile Persuasion: 20 Perspectives on the Future of Behavior Change*. Stanford Captology Media. <https://books.google.de/books?id=Ier1oHZeudQC>
- [13] Jaap Ham, René Bokhorst, Raymond Cuijpers, David van der Pol, and John-John Cabibihan. 2011. Making Robots Persuasive: The Influence of Combining Persuasive Strategies (Gazing and Gestures) by a Storytelling Robot on Its Persuasive Power. In *Social Robotics*, Bilge Mutlu, Christoph Bartneck, Jaap Ham, Vanessa Evers, and Takayuki Kanda (Eds.), Vol. 7072. Springer Berlin Heidelberg, Berlin, Heidelberg, 71–83. https://doi.org/10.1007/978-3-642-25504-5_8
- [14] Stephan Hammer. 2019. UX-orientierte Entwicklung von Empfehlungssystemen für Beratung und Assistenz. (2019).
- [15] Ryuichiro Higashinaka, Kazuki Sakai, Hiroaki Sugiyama, Hiromi Narimatsu, Tsuchihito Arimoto, Takaaki Fukutomi, Kiyooki Matsui, Yusuke Ijima, Hiroaki Ito, Shoko Araki, et al. 2017. Argumentative dialogue system based on argumentation structures. *SEMDIAL 2017 SaarDial* (2017), 154.
- [16] Anthony Hunter. 2018. Towards a framework for computational persuasion with applications in behaviour change. *Argument & Computation* 9, 1 (Jan. 2018), 15–40. <https://doi.org/10.3233/AAC-170032>
- [17] C.J. Hutto and Eric Gilbert. 2015. VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*. AAAI.
- [18] Yilin Kang, Ah-Hwee Tan, and Chunyan Miao. 2017. An Adaptive Computational Model for Personalized Persuasion. (2017), 7.
- [19] Maurits Kaptein, Joyca Lacroix, and Privender Saini. 2010. Individual Differences in Persuadability in the Health Promotion Domain. In *Persuasive Technology*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Thomas Ploug, Per Hasle, and Harri Oinas-Kukkonen (Eds.), Vol. 6137. Springer Berlin Heidelberg, Berlin, Heidelberg, 94–105. https://doi.org/10.1007/978-3-642-13226-1_11
- [20] George Konidaris, Sarah Osentoski, and Philip Thomas. 2011. Value function approximation in reinforcement learning using the Fourier basis. In *Twenty-fifth AAAI conference on artificial intelligence*.
- [21] Gernot Kräpinger. 1999. Aristoteles: Rhetorik. Übersetzt und herausgegeben von Gernot Kräpinger. Stuttgart: Reclam (1999).
- [22] Daniel J O’Keefe and Sally Jackson. 1995. Argument quality and persuasive effects: A review of current approaches. In *Argumentation and values: Proceedings of the ninth Alta conference on argumentation*. Speech Communication Association Annandale, 88–92.
- [23] Richard E Petty and John T Cacioppo. 1986. The elaboration likelihood model of persuasion. In *Communication and persuasion*. Springer, 1–24.
- [24] Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. Semeval-2015 task 12: Aspect based sentiment analysis. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*. 486–495.
- [25] Niklas Rach, Wolfgang Minker, and Stefan Ultes. 2018. Markov Games for Persuasive Dialogue. *Computational Models of Argument: Proceedings of COMMA 2018* 305 (2018), 213.
- [26] Niklas Rach, Klaus Weber, Louisa Pragst, Elisabeth André, Wolfgang Minker, and Stefan Ultes. 2018. EVA: A Multimodal Argumentative Dialogue System. In *Proceedings of the 2018 on International Conference on Multimodal Interaction - ICMI '18*. ACM Press, Boulder, CO, USA, 551–552. <https://doi.org/10.1145/3242969.3266292>
- [27] Verena Rieser and Oliver Lemon. 2011. *Reinforcement learning for adaptive dialogue systems: a data-driven methodology for dialogue management and natural language generation*. Springer Science & Business Media.
- [28] Hannes Ritschel, Tobias Baur, and Elisabeth Andre. 2017. Adapting a Robot’s linguistic style based on socially-aware reinforcement learning. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Lisbon). IEEE, 378–384. <https://doi.org/10.1109/ROMAN.2017.8172330>
- [29] Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Stefan Wagner, and Elisabeth André. 2019. Adaptive Linguistic Style for an Assistive Robotic Health Companion Based on Explicit Human Feedback. In *Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments (Rhodes, Greece) (PETRA '19)*. ACM, New York, NY, USA, 247–255. <https://doi.org/10.1145/3316782.3316791>
- [30] Ariel Rosenfeld and Sarit Kraus. 2016. Strategic Argumentative Agent for Human Persuasion.. In *ECAI*, Vol. 16. IOS Press, 320–329.
- [31] Mikey Siegel, Cynthia Breazeal, and Michael I. Norton. 2009. Persuasive Robotics: The influence of robot gender on human behavior. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, St. Louis, MO, USA, 2563–2568. <https://doi.org/10.1109/IROS.2009.5354116>
- [32] Christian Stab and Iryna Gurevych. 2014. Annotating Argument Components and Relations in Persuasive Essays. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*. Dublin City University and Association for Computational Linguistics, Dublin, Ireland, 1501–1510. <https://www.aclweb.org/anthology/C14-1142>
- [33] Richard S. Sutton, Hamid Reza Maei, Doina Precup, Shalabh Bhatnagar, David Silver, Csaba Szepesvári, and Eric Wiewiora. 2009. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*. ACM Press, Montreal, Quebec, Canada, 1–8. <https://doi.org/10.1145/1553374.1553501>
- [34] Gerben van Kleef. 2014. *Emotions as Agents of Social Influence*. Vol. 1. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199859870.013.19>
- [35] Gerben A Van Kleef, Helma van den Berg, and Marc W Heerdink. 2015. The persuasive power of emotions: Effects of emotional expressions on attitude formation and change. *Journal of Applied Psychology* 100, 4 (2015), 1124.
- [36] Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017. Computational Argumentation Quality Assessment in Natural Language. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. Association for Computational Linguistics, Valencia, Spain, 176–187. <https://www.aclweb.org/anthology/E17-1017>
- [37] Yuqiong Wang, Gale Lucas, Peter Khooshabeh, Celso De Melo, and Jonathan Gratch. 2015. Effects of emotional expressions on persuasion. *Social Influence* 10, 4 (2015), 236–249.
- [38] Klaus Weber, Hannes Ritschel, İlhan Aslan, Florian Lingenfelder, and Elisabeth André. 2018. How to Shape the Humor of a Robot - Social Behavior Adaptation Based on Reinforcement Learning. In *Proceedings of the 2018 on International Conference on Multimodal Interaction - ICMI '18*. ACM Press, Boulder, CO, USA, 154–162. <https://doi.org/10.1145/3242969.3242976>
- [39] Eric Wiewiora. 2010. Reward Shaping. In *Encyclopedia of Machine Learning*, Claude Sammut and Geoffrey I. Webb (Eds.). Springer US, Boston, MA, 863–865. https://doi.org/10.1007/978-0-387-30164-8_731
- [40] Tangming Yuan, David Moore, and Alec Grierson. 2008. A human-computer dialogue system for educational debate: A computational dialectics approach. *International Journal of Artificial Intelligence in Education* 18, 1 (2008), 3–26.