



How to Win Arguments

Empowering Virtual Agents to Improve Their Persuasiveness

Klaus Weber¹ · Niklas Rach² · Wolfgang Minker² · Elisabeth André¹

Received: 14 February 2020 / Accepted: 30 May 2020 / Published online: 16 June 2020
© The Author(s) 2020

Abstract

People make decisions every day or form an opinion based on persuasion processes, whether through advertising, planning leisure activities with friends or public speeches. Most of the time, however, subliminal persuasion processes triggered by behavioral cues (rather than the content of the message) play a far more important role than most people are aware of. To raise awareness of the different aspects of persuasion (*how* and *what*), we present a multimodal dialog system consisting of two virtual agents that use synthetic speech in a discussion setting to present pros and cons to a user on a controversial topic. The agents are able to adapt their emotions based on explicit feedback of the users to increase their perceived persuasiveness during interaction using Reinforcement Learning.

Keywords Subliminal persuasion · Adaptive agents · Persuasive dialog system · Argumentation

1 Introduction

The communication of opinions, along with different pro- and counter-arguments, is an important factor in the process of opinion building. However, people tend to get persuaded by far more than just the rational content of arguments. Presenting content-wise identical arguments in different ways, such as body language, appropriate gazing behavior as well as emotions can have a different effect on the audience's opinion towards the argument and overall stance of the topic [1, 4, 6] and, thus, the persuasive effectiveness of the conveyed content [25]. Especially in public speeches and political debates, it is not only important *what* is said (semantic content), but also *how* something is said. Looking at recent public speeches, especially by politicians, it appears that people are more likely to be influenced by the behavior or authority of speakers than by the content of the message they convey. A prominent example is the controversially

discussed *Brexit*, in which supporters have demonstrably put forward arguments with questionable validity. For instance, during the referendum campaign in 2016, it has been claimed that remaining in the EU costs around £350 million a week, which, however, disregards the fact that the UK had been granted a rebate since 1984 [7]. Although it is difficult to check what could have happened retrospectively if only logically valid arguments had been used, it still seems that people are not aware of the subliminal persuasive framing caused by their own emotions and those of the speakers.

Within this article, we present the ongoing research project EVA (Empowering Virtual Agents), in which we aim to investigate two different persuasive aspects of argumentation: (1) the effect of *what* is said and (2) the effect of *how* something is said. Thus, our approach includes logical argumentation that is focused on the content and order of the arguments (*what* to say) as well as subliminal persuasion employing non-rational argumentation, i.e. body language-based argumentation (*how* to say it). To this end, we present a novel multimodal dialog system in which two embodied virtual agents discuss a specific (controversial) topic through synthetic speech and emotions. The *what-to-say* component of the two agents is based on a logical argumentation strategy that was optimized beforehand in self-play based on objective quality criteria for effective argumentation [16]. In contrast, the persuasive power of the *how-to-say* component reflected by the agents' body

✉ Klaus Weber
klaus.weber@informatik.uni-augsburg.de

Niklas Rach
niklas.rach@uni-ulm.de

¹ Augsburg University, Augsburg, Germany

² Ulm University, Ulm, Germany

language is adapted during the interaction with the user. In particular, the user may provide explicit feedback on whether he or she finds the presentation convincing, not convincing or has a neutral opinion about it. By allowing users to tune the *how-to-say* component of the agents, we enable them to explore the impact of subliminal cues on the persuasive power of arguments in order to raise awareness that people get influenced by adapted behavior.

In Sect. 2, we give an overview of the research of persuasion. Sect. 3 describes the overall approach, a new formalization of the merged learning approach and a thorough overview of the currently employed logical and non-logical policies, while Sect. 4 shows the results at the current stage of the research.

2 Research of Persuasion

2.1 General Theory

The theory of persuasion goes back to the Greek philosopher Aristotle, who identified three means of persuasion, that are (1) *logos* (the logical and rational aspects, i.e., the content), (2) *pathos* (the emotional engagement between persuader and persuadee) and (3) *ethos* (the persuader's personality and character) [11].

Psychological persuasion models developed by Petty et al. [12] (Elaboration Likelihood Model – ELM) and Chaiken et al. [3] (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 [3]. Instead, 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*” [12].

2.2 Non-verbal Signals and Persuasion

The persuasive effect of non-verbal signals in Human-Agent-Interaction has extensively been investigated. For instance, it has been shown that robots using non-verbal cues (body language, gaze, ...) are perceived as more persuasive than the ones not using them. Examples are given by Chidambaram et al. [4], who compared the effect of non-verbal cues (gesture, gaze, ...) to vocal cues showing a higher persuasive effect for non-verbal cues. Further,

Ham et al. [6] observed that gazing increased the persuasive effect more than using gestures alone did. Also, it has been found by Andrist et al. [1] that practical knowledge and rhetorical ability can affect the persuasive effect. The EASI theory (Emotion As Source of Information, Kleef et al. [23]) states that both intra- and interpersonal influences of emotions exist, i.e. the persuadee's own emotions can affect the outcome of a persuasive message as well as the emotions of the persuader conveying the persuasive message. According to this theory, two different processes exist, namely *inference* and *affective reaction*. In addition to that, two moderator classes determine which of these two processes takes precedence: *appropriateness* and *information processing*. The latter one depends, similar to the ELM, on the motivation and ability of the persuadee to engage in thorough information processing (*epistemic motivation*). Studies have proven this theory and shown that people use the emotions of the source as an information channel when forming their own attitudes [24]. In accordance with this, DeSteno et al. [5] showed that persuasive messages are more successful if they are framed with emotional overtones that correspond to the emotional state of the recipient.

The whole argumentation of Brexit shows that the people's opinion is strongly driven by emotions and not just by the validity of arguments. There is also 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 subliminal persuasion.

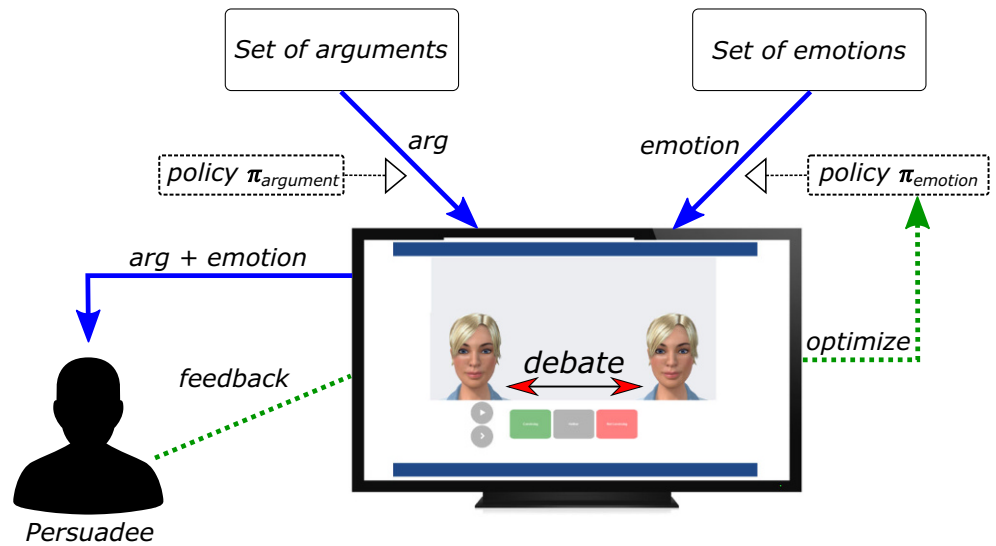
3 Concept and Approach

In the following, we present an overview of our approach and provide a formal description of the overall concept at a high level (see Sect. 3.1) as well as a thorough discussion of the two employed policies: (1) Policy $\pi_{emotion}$ for emotional behavior generation (see Sect. 3.2.2) and (2) policy $\pi_{argument}$ for the argumentation strategy (see Sect. 3.2.1).

For every interaction step t , each agent selects an utterance (one after the other) as well as a corresponding emotion and presents it to the user based on its learned *merged policy* (see Def. 4). The user provides the agents with feedback, which is used to optimize the policy (here only $\pi_{emotion}$, see Fig. 1¹) based on Reinforcement Learning (RL).

¹ Freely available for non-profit academic research <https://charamel.com/>.

Fig. 1 Overview of our approach with two virtual agents¹ debating about a controversial topic employing a pre-trained logical strategy π_{argument} as well a dynamically learned behavior strategy π_{emotion} adapted to the user/persuadee based on user's explicit feedback how convincing an argument is perceived (convincing, neutral, not convincing). This system is an extension of the rule-based approach presented in [18] (including trained logical strategy) and the adaptation approach (without trained logical strategy) presented in [28]



3.1 Formalization

Formally, our approach can be described as Markov Game with i players (here: 2) as 5-tuple $(\mathcal{I}, \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T})$, where \mathcal{I} defines the set of players, \mathcal{S} defines the merged state space, \mathcal{A} the joint action space and \mathcal{R} the joint reward function and $\mathcal{T} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0,1]$ determines the transition function as follows [2]:

Definition 1 (Merged State Space) Let $\mathcal{S}_{\text{emotion}}$ be the sub state space in order to determine the next emotion and $\mathcal{S}_{\text{argument}}$ the sub state space for determining the next argument, then the system's merged state space \mathcal{S} is defined as

$$\mathcal{S} := \mathcal{S}_{\text{argument}} \times \mathcal{S}_{\text{emotion}} \quad (1)$$

Definition 2 (Joint Action Space) Let $\mathcal{A}_{\text{emotion},i}$ be the emotion sub action space and $\mathcal{A}_{\text{argument},i}$ the argumentation sub action space for player $i \in \mathcal{I}$. Further, let $\mathcal{A}_i := \mathcal{A}_{\text{emotion},i} \times \mathcal{A}_{\text{argument},i}$ be the respective merged action space for player $i \in \mathcal{I}$, then the joint action space is defined as

$$\mathcal{A} := \times_{i \in \mathcal{I}} \mathcal{A}_i \quad (2)$$

Definition 3 (Joint Reward Function) Let the function $\mathcal{R}_{\text{emotion},i} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ be the reward given for the emotional part and $\mathcal{R}_{\text{argument},i} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ be the reward given for the rational part for player $i \in \mathcal{I}$.

Further, let $\mathcal{R}_i := \mathcal{R}_{\text{emotion},i} \times \mathcal{R}_{\text{argument},i}$ be the respective merged reward function, then the joint reward function \mathcal{R} is defined as

$$\mathcal{R} := \times_{i \in \mathcal{I}} \mathcal{R}_i \quad (3)$$

Definition 4 (Joint Policy) Let $\pi_{\text{emotion},i}$ be the emotion policy and $\pi_{\text{argument},i}$ be the argumentation policy for player $i \in \mathcal{I}$. Further, let $\pi_i := \pi_{\text{emotion},i} \times \pi_{\text{argument},i}$ be the respective merged policy with $\pi_i : \mathcal{S} \rightarrow \mathcal{D}(\mathcal{A}_i)$ the distribution of possible actions given a state $s \in \mathcal{S}$, then the joint policy is defined as

$$\pi := \times_{i \in \mathcal{I}} \pi_i \quad (4)$$

In the herein presented Markov Game, the agents select and execute their actions one after the other (and not contemporaneously). Thus, we define an awaiting action $a_w \in \mathcal{A}_i$ for every agent $i \in \mathcal{I}$ who is not in turn. Further we denote any state $s_t \in \mathcal{S}$ (respectively action $a_t \in \mathcal{A}_i$) for agent $i \in \mathcal{I}$ at time step t as s_t^i (a_t^i). It is further notable that both agents have opposing goals, which is formally called a zero-sum game. The optimum joint policy π corresponds to a Nash Equilibrium, which means that a change of any strategy π_i does not yield an advantage for the respective agent if the policy of the opposing player is kept stationary. It is important to note that the two state and action spaces clearly differ in their size and complexity. Since the argumentation policy $\pi_{\text{argument},i}$ should be able to differentiate between all available arguments, a real-time optimization (during the interaction with a user) is not feasible. On the other hand, research has yielded different objective quality criteria for logical argumentation which

can be utilized in order to define an appropriate reward. Since this is not necessarily the case for the emotional policy $\pi_{emotion,i}$ (due to subjectivity), it has to be adapted in real-time and directly to the user response. Hence, we split the learning process into two phases (pre-training and real-time) and optimize the two parts of the policy π separately at the current state of our work. Both learning phases are discussed in more detail in the following subsections. The Algorithm 1 sketches the general procedure during the interaction.

Algorithm 1: Adaptation

```

foreach  $t = 1, \dots, n$  do
  foreach  $i \in \mathcal{I}, 1 \leq i \leq m$  do
     $s_t^i \leftarrow (s_{argument,t}^i, s_{emotion,t}^i) \in \mathcal{S}$ 
    select action  $a_{argument,t}^i$  for  $s_{argument,t}^i$ 
    select action  $a_{emotion,t}^i$  for  $s_{emotion,t}^i$ 
    apply  $a_t^i = (a_{argument,t}^i, a_{emotion,t}^i)$ 
    observe temporary state  $s_{t+1}^i$ 
    get reward
     $\mathcal{R}_i(s_t^i, a_w^0 \times \dots \times a_t^i \times \dots \times a_w^m, s_{t+1}^i)$ 
    update policy  $\pi_{emotion,i}$ .
  end
end
  
```

3.2 The Policy π

Since persuasion does not only depend on the content but also the way, how a message is conveyed, the *joint policy* π employed by the agents contains two merged sub-policies, that are $\pi_{argument}$ and $\pi_{emotion}$ as formally (high-level) described in Sect. 3.1. In the following section, we give a more in-depth overview of the two policies separately.

3.2.1 Learning of the Argumentation policy $\pi_{argument}$

A significant amount of research has been conducted in order to formalize arguments and argumentation strategies in dialogues (see for example [15] for an overview). Our approach builds on these results and utilizes a dialog game framework tailored to persuasion in order to structure the interaction between the two agents. The respective dialog game can formally be described as tuple (\mathcal{L}, D) , with \mathcal{L} logic for defeasible argumentation and D the dialog system proper [14]. Consequently, \mathcal{L} encodes a formal representation of the available arguments (including their relations to each other), and D defines the rules for the interaction. This includes aspects like turn-taking, winning criteria and allowed moves in each state of the interaction.

Within this project, we utilize the dialog game for argumentation introduced in [13], since it was motivated by providing a flexible framework that still ensures consistent dialogs. The framework allows for five different speech acts which are *claim*(ϕ_i), *argue*(ϕ_i , so ϕ_j), *why*(ϕ_i), *retract*(ϕ_i)

and *concede*(ϕ_i) with ϕ_i and ϕ_j argument components in \mathcal{L} . The available argument components $\{\phi_i\}$ are encoded following the argument annotation scheme introduced in [21]. This choice is motivated by the ambition to combine the presented system with existing argument mining approaches to ensure flexibility of the system in view of the discussed topics. The annotation scheme allows for three different types of components (*Major Claim*, *Claim*, *Premise*). It is important to note that a *Claim* component in the argument structure is not related to the *claim* speech act in the dialog game since they are both part of separate frameworks. In addition to the three types of components, the argument annotation scheme also allows for two relations between components that are *support* and *attack*. Each component apart from the *Major Claim* (which has no relation) has exactly one unique relation to another component. We denote the set of arguments that can be constructed from the annotated set as *Args*. The arguments $arg_i \in \text{Args}$ have the form $arg_i = (\phi_i, \text{so } \phi_j)$ if ϕ_i supports ϕ_j and $arg_i = (\phi_i, \text{so } \neg\phi_j)$ if ϕ_i attacks ϕ_j . Each argument $arg \in \text{Args}$ refers to one of two existing *stances* $\in \{+, -\}$ of the topic. For the sake of simplicity, the *Major Claim* is defined as $arg_0 = \phi_0$. The complete setup including a discussion of an annotated structure and a preliminary study on how the resulting artificial argumentation was perceived by a human audience was presented in [17].

In the next step, the aim is to explore the freedom within the dialog framework to optimize the logical strategy of the agent. In order to do so, we focus again on a formalization of the problem as a Markov Game which allows addressing policy optimization as a multi-agent Reinforcement Learning task. The advantage of the Markov Game formalism is that it does not depend on a pre-defined strategy or additional annotated training data, assuming that a formal reward function is given. Moreover, it allows the agent to explore different (and maybe unknown) strategies instead of imitating a human policy. The formal description can be derived from the general formalism in Sect. 3.1 by excluding the emotional aspects from the formal description. This is possible since the introduced formalization assumes a separation of the two aspects in the action and state space as well as in the reward function and the policy. The utilized formal translation of a generalized dialog game for argumentation into a Markov Game was extensively discussed in [16]. The reward functions $\mathcal{R}_{argument,i}$ of both agents $i \in \mathcal{I}$ are currently based on the winning criterion of the underlying dialog game for argumentation since it provides a formal indicator for the outcome of the interaction. However, since also the perception of the content is highly subjective, our ongoing work is focused on including user feedback into the reward and combining it with the formal aspects. The complete pipeline, including two argument structures, the

dialog game for argumentation, optimized policy $\pi_{argument}$ and a virtual agent were presented in [18].

3.2.2 Learning of the Emotion policy $\pi_{emotion}$

The sub policy $\pi_{emotion}$ is – in contrast to the argumentation policy $\pi_{argument}$ – learned during interaction with the user. In order to enable the agents to adapt their behavior, we employ Reinforcement Learning as it is a suitable tool for real-time adaptation in human-agent scenarios as shown in [20, 26, 27].

The policy $\pi_{emotion}$ depends on the argumentation policy $\pi_{argument}$ since the type of argument (or at least the content) defines what kind of emotion is the right one to be chosen. Consequently, the sub-state $s_{emotion}$ is derived from the argumentation sub action $a_{argument} \in \mathcal{A}_{argument}$ as described in the following.

Definition 5 (State Space $\mathcal{S}_{emotion}$) Let $a_{argument,t}^i = (arg) \in \mathcal{A}_{argument,i}$ be the sub argument action with $arg \in Args$ chosen at time step t for agent $i \in \mathcal{I}$, $stan : Args \rightarrow \{+, -\}$ the stance and $rel : Args \rightarrow \{attack, support\}$ the relation of argument arg . Further let $score : Args \rightarrow [-1, 1]$ be the normalized compound score of the sentiment analysis of an argument. Then the sub state $s_{emotion}$ is defined as follows considering it is player i 's turn:

$$s_{emotion} := (i, stan(arg), rel(arg), score(arg)) \quad (5)$$

Definition 6 (Action Space $\mathcal{A}_{emotion}$) The discrete action space $\mathcal{A}_{emotion}$ consists of employed emotions, such as *happy*, *disappointed*, *sad*, *angry*, which can be both displayed by facial expressions as well as gestures.

To this end, we employ the user's feedback (*convincing*, *neutral*, *not convincing*) to compute a prediction of the current user's stance during interaction as introduced in [28]. The employed *prediction model* is based on *bipolar-weighted argument graphs* (BWAG) that assigns a weight $w_i \in [0, 1]$ to each argument $arg_i \in Args$, which is used for computing the argument's strength s_i considering the strengths of its child arguments in the argument graph (see Fig. 2).

The strength of the root argument s_0 is finally taken as the predicted user's stance ρ_t (Def. 7) at time step t . This approach has the advantage that the underlying argument structure is taken into account for the learning process [28]. Consequently, the reward function $\mathcal{R}_{emotion}$ is defined as:

Definition 7 (Reward Function $\mathcal{R}_{emotion}$) Let $f_t : Args \rightarrow \{0.0, 0.5, 1.0\}$ be the user feedback function defining how convincing an argument is perceived by the user at time step t – ranging from 0 (*not convincing*) to 1 (*convincing*).




Fig. 2 Sketch of a BWAG for an arbitrary argument $m \in Args$ having two children k and $k+1$. The strength s_m for argument m is computed by a function g considering its own weight w_m and the strength of the child nodes s_k and s_{k+1} (taken from [28])

Further let $\rho_t^i : (f_1, \dots, f_t) \rightarrow [0, 1]$ a function that maps the feedback signals to a prediction of the current user's stance until time step t (see Weber et al. [28] for more information about the prediction model) with respect to agent $i \in \mathcal{I}$, then the reward $\mathcal{R}_{emotion,i,t}$ at time step t is defined as:

$$\mathcal{R}_{emotion,i,t} := \rho_t^i(f_1, \dots, f_t) - \rho_{t-1}^i(f_1, \dots, f_{t-1}) \quad (6)$$

with $\rho_0 = 0.5$

While employing the user's feedback directly as a reward signal allows for quick adaptation, using a prediction of the user's stance allows for learning a more fine-grained strategy depending on the actual outcome of the debate. This is because the prediction model allows the agents to observe their current position with respect to the opponent. The advantage of that kind of approach is that an agent can learn to take over the opponent's strategy when it notices that it currently holds the worse position. Further, argument-specific and structure-specific information can be used to provide more fine-grained information for learning. In addition to that, behavior learning can be combined with fine-grained logical strategies [28].

In the next step, we aim to explore the effect when adapting emotions to the listener in order to both verify that an adaptation leads to higher persuasive effectiveness and to raise awareness of the subliminal influence of adapted emotions.

3.3 Dialog Example

To give the reader a better understanding of the overall system, we sketch in the following a detailed example of an interaction between both agents and the user.

The herein presented dialog between the two agents concerns the claim $\phi_0 = \text{Marriage is an outdated institution}^2$. The stances are as follows:

- (i) Left agent: $\phi_0 \rightarrow \text{stan} = +$
- (ii) Right agent: $\neg\phi_0 \rightarrow \text{stan} = -$

For the sake of simplicity, we assumed that the user is against the claim ($\neg\phi_0$) and finds the right agent (*agent 2*) convincing, if attacking arguments are presented with an *angry* emotion. The left agent (*agent 1*) is found convincing, if the sentiment of the argument is negative. In all other cases, a negative reward is given. Fig. 3 shows the example dialog between the two agents.

The first argument given by *agent 1* is presented with a neutral emotion and, therefore, *not convincing* leading to a negative reward for the agent. The next argument of *agent 2* is an attacking argument and presented along with an *angry* emotion and, therefore, found *convincing*. *Why* speech acts do not require feedback since they do not contain any argument, therefore, the reward is always zero. The fourth argument is found *not convincing* due to the nature of our fictitious user, while the last argument given by *agent 1* is found *convincing* again since the correct behavior is chosen.

4 Evaluation Results

In the following, we highlight the results of our research at the current stage. As we split the learning process into two phases, we discuss each of them separately.

Argumentation policy. For the argumentation strategy, the combination of argument structure and dialogue game in a multi-agent setup was evaluated in a user study [17] by comparing transcripts of artificial discussions with human-generated ones. In the course of the survey, each participant answered questions regarding the logical consistency of the dialogs, the argumentation strategy of the agents and the naturalness of the interaction on a five-point Likert Scale. The results showed that the agent-agent dialogues were perceived as logically consistent but also a significantly higher perceived naturalness of the interaction for human-generated dialogues. We concluded that the utilized methods are generally suitable for our task (as they lead to cogent interactions) but that the perceived low naturalness indicates room for improvement in view of the perceived quality of the argumentation.

In the next step, the Markov Game formalism was applied in order to explore the freedom of choices for the

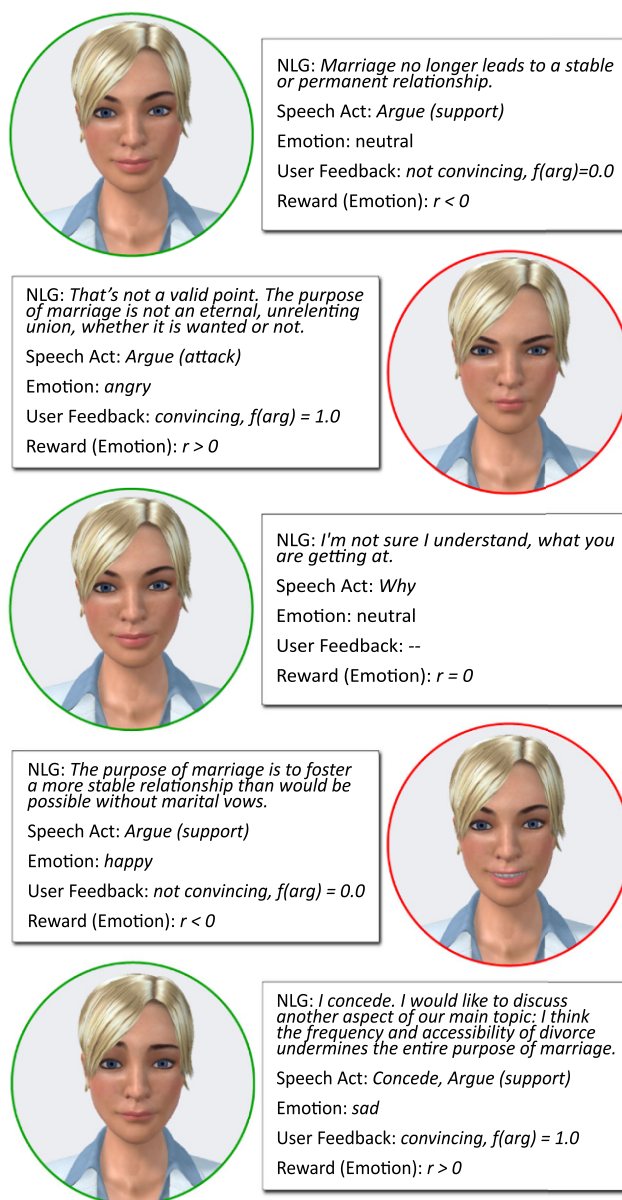



Fig. 3 Example dialog between the two agents showing different arguments presented with emotions and the given feedback of an exemplary fictitious human user

agents that is provided by the utilized dialogue game framework. We evaluated the approach in a proof of principle setup with a reward function for which the optimal policy is known [16]. We trained a policy for both sides of the discussion by means of RL for ten randomly generated argument structures in order to exclude topic dependencies of the results. The trained policies were evaluated against the known optimal policy based on probabilistic rules and shown to lead to the expected outcome. We concluded that for the investigated case, the optimal policy could be found by means of the applied techniques.

² Material reproduced from www.iedebate.org with the permission of the International Debating Education Association. Copyright © 2005 International Debate Education Association. All Rights Reserved.

Fig. 4 Simulation results including 95% confidence intervals depicting the cumulative prediction ρ_t showing a continuous increase over time even for high noise (taken from [28])



Emotion policy. We have evaluated the adaptation of the emotional policy in a simulative setup showing that the agents are able to increase their perceived persuasive effectiveness (ρ_t , Fig. 4) by means of using the provided human feedback signals even for high noise (as user feedback is not deterministic) [28].

Further, we evaluated the prediction model in a user study with 48 participants in order to verify its practical potential, accuracy, and validity. For this purpose, the participants were asked to listen to an agent, who brought pro and counter arguments for visiting a hotel, and asked to give direct feedback whether they found an argument convincing. This feedback was used to calculate the prediction ρ_t of the user's stance (see Definition 7). In a post-study questionnaire, the participants were asked whether or not they like to visit the hotel. We then evaluated to what degree the predicted user's stance and the subjective decision match. Table 1 and Fig. 5 summarize the results [28], which show that the prediction model of the user's stance is very accurate making it a suitable measure for rewarding the agents during the interaction, especially in the herein presented ap-

proach because the agents can use this information to adopt strategies of the opponent if they have a worse position in the discussion.

5 Discussion of Limitations

There are two (potential) limitations that need to be discussed for the sake of completeness.

Argument content not considered in the emotion policy. In Sect. 3.2.2, we presented the adaptation approach of the emotion policy $\pi_{emotion}$. The RL state space (Definition 5) was defined by a triple consisting of the *stance*, the *relation* and the *sentiment score* of the selected argument. So, at the current stage of this work, the emotional policy is quite independently learned from the argumentation policy and the content of the argument is not considered at all. However, there might be an interaction between the content of the argument and the expressed emotion. The issue, however, is that real-time adaptation requires a small state space [27]. Therefore, we limited the state space to three dimensions. To include at least some content-related information, we added the *sentiment score*, which, however, just describes the negativity of the argument and not the direct content. Assigning a single score to arguments describing the content is to our best knowledge not possible. Therefore, it would be necessary to include all arguments in the state space so that the agent learns a policy that takes the content of the argument into account. This would exponentially increase the state space and is therefore not suitable for real-time adaptation. Another possible solution could be to add an assessment score describing the quality of the argument by employing automatic quality assessment techniques of arguments as recently proposed by Toledo et al. [22]. In summary, additional research and studies are needed to overcome this limitation or to show that the proposed state space contains sufficient information to learn a consistent emotional policy successfully.

Table 1 F1 score for different prediction confidences

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




Fig. 5 Model accuracy of predicted user's stance depending on different confidence values

Perception of arguments depends on the agent's focus.

Within the experiment to validate the prediction model, all arguments were presented by a single agent in a user-directed manner (uni-directional persuasion [8]). However, there is evidence from the literature that multiple agents overall appear to be more persuasive compared to single agents in case of maintaining a consistent argumentation strategy. This effect is notably increased by using vicarious bi-directional persuasive agents [9] (in which one agent tries to persuade the other agents and indirectly persuades the observer of the dialog [8]). Consequently, one might wonder whether the evaluated prediction model performs equally well within the proposed multi-agent scenario since the persuasive outcome could be different. However, we argue that since our model depends only on direct user feedback and since each argument (not just the whole debate) would be perceived differently in the first place, the user feedback received would be different, and thus the predicted stance of our model would implicitly take such effects into account. Nevertheless, it would be worth conducting a follow-up study to verify this assumption.

6 Summary and Outlook

Within this work, we have provided an overview of the EVA project that aims for synthesizing and combining different aspects of argumentation. Our proposed system is comprised of two interacting agents that are represented by virtual avatars. The argumentative aspects, addressed in this paper, cover (1) the learning of an optimal logical strategy within a dialog game framework from objective optimization criteria and (2) a real-time adaptation of the emotional presentation of arguments based on explicit user feedback and a computed prediction of the current user's stance over time. Our results showed the predictive power of the proposed prediction model and that the trained logical policy performs comparably well as the optimal policy based on probabilistic rules.

In our future work, we will explore different techniques to implicitly estimate the persuasive effect of the presented arguments on the user without explicit feedback based on the approaches discussed in [19] and the prediction model presented in [28]. Since argumentation as a whole is highly subjective [10], we aim to extend the system to allow the agents to adapt both investigated aspects (*how* and *what*) to the user simultaneously in order to determine the most effective combination of the different aspects. Thus, we will include a combination and simultaneous learning of both aspects of the strategy. To this end, we will explore different approaches, including function approximation and fine-tuning of pre-trained strategies. Finally, we aim to com-

pare the respective outcomes with the ones achieved by the current separated approach.

Acknowledgements 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).

Funding Open Access funding provided by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Andrist S, Spannan E, Mutlu B (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, pp 341–348
2. Barlier M, Perolat J, Laroche R, Pietquin O (2015) Human-machine dialog as a stochastic game. In: Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and dialog, pp 2–11
3. Chaiken S, Liberman A, Eagly AH (1989) Heuristic and systematic information processing within and beyond the persuasion context. In: Uleman JS, Bargh JA (eds) Unintended Thought. New York, Guilford, pp 212–252
4. Chidambaram V, Chiang Y-H, Mutlu B (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, Boston, pp 293–300
5. DeSteno D, Petty RE, Rucker DD, Wegener DT, Braverman J (2004) Discrete emotions and persuasion: the role of emotion-induced expectancies. *J Pers Soc Psychol* 86(1):43
6. Ham J, Bokhorst R, Cuijpers R, van der Pol D, Cabibihan J-J (2011) Making robots persuasive: the influence of combining persuasive strategies (gazing and gestures) by a storytelling robot on its persuasive power. In: Mutlu B et al (ed) Social Robotics. Springer, Berlin, Heidelberg, pp 71–83
7. Henley J (2016) Why Vote Leave's £350m weekly EU cost claim is wrong. <https://www.theguardian.com/politics/reality-check/2016/may/23/does-the-eu-really-cost-the-uk-350m-a-week>. Accessed 18 Dec 2019 (The Guardian News and Media)
8. Hunter A (2018) Towards a framework for computational persuasion with applications in behaviour change. *Argum Comput* 9(1):15–40
9. Kantharaju RB, De Franco D, Pease A, Pelachaud C (2018) Is two better than one? Effects of multiple agents on user persuasion. In: Proceedings of the 18th International Conference on Intelligent Virtual Agents, pp 255–262

10. O’Keefe DJ, Jackson S (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, pp 88–92
11. Krapinger G (1999) *Aristoteles: Rhetorik*. Reclam, Stuttgart (Translated and published by Gernot Krapinger)
12. Petty RE, Cacioppo JT (1986) The elaboration likelihood model of persuasion. In: *Communication and persuasion*. Springer, New York, pp 1–24
13. Prakken H (2000) On dialogue systems with speech acts, arguments, and counterarguments. In: Ojeda-Aciego M, de Guzmán IP, Brewka G, Moniz Pereira L (eds) *Logics in artificial intelligence. JELIA 2000. Lecture notes in computer science*, vol 1919. Springer, Berlin, Heidelberg
14. Prakken H (2005) Coherence and flexibility in dialogue games for argumentation. *J Log Comput* 15(6):1009–1040
15. Prakken H (2018) Historical overview of formal argumentation vol 1. College Publications, London, pp 73–141
16. Rach N, Minker W, Ultes S (2018) Markov games for persuasive dialogue. In: *Proceedings of the 7th International Conference on Computational Models of Argument* Warsaw, September
17. Rach N, Langhammer S, Minker W, Ultes S (2019) Utilizing argument mining techniques for argumentative dialog systems. In: *9th International Workshop on Spoken Dialog System Technology*. Springer, Singapore
18. Rach N, Weber K, Pragst L, André E, Minker W, Ultes S (2018on) EVA: a Multimodal argumentative dialogue system. In: *Proceedings of the 2018 on International Conference on Multimodal Interaction ICMI ’18*. ACM, Boulder, CO, pp 551–552
19. Rach N, Weber K, Aicher A, Lingenfelser F, André E, Minker W (2019) Emotion recognition based preference modelling in argumentative dialog systems. In: *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, Kyoto, pp 838–843
20. Ritschel H, Baur T, André E (2017) Adapting a robot’s linguistic style based on socially-aware reinforcement learning. In: *26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, Lisbon, pp 378–384
21. Stab C, Gurevych I (2014) Annotating argument components and relations in persuasive essays. In: *COLING*, pp 1501–1510
22. Toledo A, Gretz S, Cohen-Karlik E, Friedman R, Venezian E (2019) Automatic argument quality assessment – new datasets and methods. In: *EMNLP 2019*. arXiv:1909.01007
23. van Kleef G (2014) Emotions as agents of social influence. In: Harkins SG, Williams KD, Burger J (eds) *The Oxford Handbook of Social Influence*. Oxford University Press, Oxford
24. Van Kleef GA, van den Berg H, Heerdink MW (2015) The persuasive power of emotions: effects of emotional expressions on attitude formation and change. *J Appl Psychol* 100(4):1124
25. Wachsmuth H, Naderi N, Hou Y, Bilu Y, Prabhakaran V, Alberdingk Thijm T, Hirst G, Stein B (2017) Computational argumentation quality assessment in natural language. In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pp 176–187
26. Weber K, Ritschel H, Lingenfelser F, André E (2018) Real-time adaptation of a robotic joke teller based on human social signals. In: *Proceedings of the 17th International Conference on Autonomous Agents and Multi Agent Systems AAMAS ’18*, Richland, SC, pp 2259–2261
27. Weber K, Ritschel H, Aslan I, Lingenfelser F, André E (2018) How to shape the humor of a robot – social behavior adaptation based on reinforcement learning. In: *Proceedings of the 20th International Conference on Multimodal Interaction ICMI ’18*. ACM, Boulder, CO, pp 154–162
28. Weber K, Janowski K, Rach N, Weitz K, Minker W, Ultes S, André E (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*, Auckland, NZ, pp 1476–1484