Pre-print of full paper presented at CONVERSATIONS 2017 - an international workshop on chatbot research and design, November 22, in conjunction with the 4th International Conference on Internet Science, Thessaloniki, Greece. The final version of the paper will be published in the conference post-proceedings as part of Springer LNCS.

# Creating Dialogues using Argumentation and Social Practices

Frank Dignum and Floris Bex

Department of Information and Computing Sciences, Utrecht University, The Netherlands {F.P.M.Dignum,F.J.Bex}@uu.nl

Abstract. Research in chatbots is already more than fifty years old, starting with the famous Eliza example. Although current chatbots might perform better, overall, than Eliza the basic principles used have not evolved that much. Recent advances are made through the use of massive learning on huge amounts of resources available through Internet dialogues. However, in most domains these huge corpora are not available. Another limitation is that most research is done on chatbots that are used for focused task driven dialogues. This context gives a natural focus for the dialogue and facilitates the use of simple reactive rules or frame-based approaches. In this paper, we argue that if chatbots are used in more general domains we have to make use of different types of knowledge to successfully guide the chatbot through the dialogue. We propose the use of argumentation theory and social practices as two general applicable sources of knowledge to guide conversations.

# **1** INTRODUCTION

With the current surge in AI the interest in using chatbots, for example, for client contacts, has also increased. Due to the amazing results of deep learning and the excellent framing of the capabilities of their chatbots by companies like Apple and Google, people expect chatbots to "learn" to give the "correct" reaction to any user input. However, whereas techniques from deep learning can be used in some particular domains or applications, they are of no use in other domains. For example, learning the right behaviour becomes difficult when only small amounts of data is available. Furthermore, some features of interactions are of a social nature and can only be inferred in an indirect way using knowledge of the context, which is not represented in the textual data.

Chatbots have mainly been successful in situations where the user interacts with the chatbot for a very specific purpose that is determined beforehand. For example, in a recent system [4] we use a chatbot to interact with a user that is reporting fraud about an e-commerce transaction. Even though this application is far from simple it is clear from the start which information is needed from the user and what the goal of the interaction is. The user also sees the form and, even though the interaction is in natural language, the context is very focused and clear. For example, when the system asks what happened, it is clear that only events pertaining to the particular e-commerce transaction should be related. Thus the parsing of the input can concentrate on keywords and particular phrases indicating order of events.

More interesting (and difficult) situations arise when the purpose of an interaction is to find out what exactly the needs of the user are. This can be e.g. a user finding out whether it is interesting for him to change his mortgage to another format, or whether a user would need a license to extend the roof of his house with a standard dormer structure or whether insisting on a crime report might lead to a possible restitution of the damages. In all these cases the purpose of the interaction is no longer to determine parameters for a specific action of the organization, but rather it is the joint determining which type of action would be appropriate. The main difference seems to be that the dialogue is becoming a real two-way dialogue where both parties seek and give information and are planning steps in the dialogue which are not fixed beforehand.

We argue that for the above situations we should take an agent-based approach. Current chatbot technology is directed to create a response that makes sense in a particular context. Thus there are means to keep track of the dialogue state and adapt responses based on history of the dialogue. However, the technology is not geared towards creating the dialogue structure in collaboration with the user. Agent technology gives a different perspective on the dialogue, where the agent is not only reacting to input of the user, but also can pursue its own goals or joint goals with the user. This leads to dialogues as joint intentions in the sense of [7]. Taking this perspective also facilitates two approaches that support dialogue management. First of all it facilitates the use of argumentation in dialogues [13]. The agent can reason about the input of the client as being an argument in the context of the dialogue history. This position indicates the type of reaction the agent can give and for what purpose the reaction is intended. Argumentation can also be used to structurally guide the dialogue to a conclusion accepted by both parties.

Secondly, we can use social practices as concepts that shape our social interactions. Social practices describe many aspects of the context of the interaction and thus give rise to expectations of what will happen at what stage of the interaction, who has initiative, etc. For example, when a client of a bank contacts the bank he can expect that the bank wants to know his bank account number. This will facilitate further interaction because through this number the bank can retrieve all information about the client. Also in many formal interactions the organization will summarize and check all agreements that have been made with the client at the end of the meeting. These types of dialogue patterns not only give a context to parts of a dialogue but help shape the interaction and focus the attention and direct actions towards a specific joint goal to be achieved. These aspects can be used to modularize the rules of interaction and focus attention to topics fitting in each phase of a dialogue. This will limit the freedom of the natural language used in a very natural way and also in many case convert natural language parsing into pattern matching with the pattern being the expected input at a certain time in the dialogue.

In the rest of this paper we will give some background on social practices and argumentation in section 2. Subsequently we will show in section 3 how these can be specifically useful for chatbots. In section 4 we give some conclusions and pointers to future work.

# 2 Background

### 2.1 Social Practices

Social practices can be seen as patterns which can be filled in by a multitude of single and often unique actions [17], that endure between and across specific moments of enactment. Through (joint) performance, the patterns provided by the practice are filled out and reproduced. Each time it is used, elements of the practice, including know-how, meanings and purposes, are reconfigured and adapted [18]. Therefore the use of social practices includes a constant learning of the individuals using the social practice in ever changing contexts. In this way social practices guide the learning process of agents in a natural way.

In [18] the social aspect of social practices is emphasized by giving the social practice center stage in interactions and letting individuals be supporters of the social practice. It shows that social practices are shared (social) concepts. The mere fact that they are shared and jointly created and maintained means that individuals playing a role in a social practice will expect certain behaviour and reactions of the other participants in the social practice. Thus it is this aspect that makes the social practices so suitable for use in designing interactive systems and especially where the interaction is not completely structured beforehand such as in dialogues between the system and the user.

**Characteristics of Social Practices** Researchers in social science have identified three broad categories of elements of practices [10]:

- Material: covers all physical aspects of the performance of a practice, including the human body (relates to physical aspects of a situation).
- Meaning: refers to the issues which are considered to be relevant with respect to that material, i.e. understandings, beliefs and emotions (relates to social aspects of a situation)
- Competence: refers to skills and knowledge which are required to perform the practice (relates to the notion of deliberation about a situation).

These components are combined by individuals when carrying out a practice. Each individual embeds and evolves (through conditioning) meaning and competence, and adopts material according to its motives, identities, capabilities, emotions, and so forth, such that it implements a practice. Individuals and societies typically evolve a collection of practices over time that can be applied in different situations. Although social practices provide a handle for modelling the deliberation of socially intelligent agents because they seem to combine the elements that we require for socially intelligent behaviour, they are a relatively novel and vaguely defined concept from sociology that cannot be just applied in agent systems. Thus we have tried to convert the above very general and vague definition into a slightly more concrete and usable format. Table 1 informally defines the aspects that are relevant in a typical social practice.

The material aspects of social practices are divided over the (physical) context, activities and plan patterns. The physical context contains elements from the environment where we distinguish between objects (resources) and actors and the locations of all elements. Location can again be physical location but also temporal location (thus the time of day or year). The activities indicate the normal activities that are expected within the practice. In a doctor consult these are greeting, explaining symptoms, discussing treatment, recapture findings, questioning information, etc. Note that not all activities need to be performed. They are meant as potential courses of action. In the same way the plan patterns describe usual patterns of actions that can be used as frames from which a planning process can start.

The meaning of a social practice is divided over the social interpretation of the context and the social meaning of the activities performed in the practice. The first part actually determines the social context in which the practice is used. Of course, this part can be highly subjective with different agents interpreting signals in their own way. The social context also contains the roles and norms that are attached to the situation. E.g. a person can be a doctor or patient, which leads to different expectations of this person. Norms also give rules of (expected) behaviour within the practice. The part that we have now coined the meaning of the practice refers to the social meaning of the activities that are (or can be) performed in the practice. Thus they indicate social effects of actions. The competencies are still present in the form of capabilities the agent should have to perform the activities within this practice.

The table describing the social practices as given above starts resembling the list of aspects that also play a role in agent organization models (see e.g. [8]). It can be seen as an analogue of the connection between imposed and emerging norms. Both organizations and social practices give a kind of structure to the interactions between agents. However, organizations provide an imposed (top-down) structure, while the social practices form a structure that arises from the bottom up. Thus where organizational interaction patterns indicate minimal patterns that agents should comply with, the patterns in a social practice indicate minimal patterns that can and are usually used by the agents.

A final, very important, observation is that social practices are not the same as scripts. They do not generate a complete, deterministic protocol, but rather give a set of expectations that can be used to guide the interaction, but also can be deviated from. Secondly, social practices do not just indicate expected actions, but also give information about expected resources, information and timing. Moreover, they give social meaning to actions in the context of the social practice that facilitate social reasoning and planning during the interaction.

	Meaning
Context	
Actors	The actors in the social practice.
Roles	The roles such as client, consultant, account manager
Resources	The set of resources available for the social practice. This can be icons
	on a joint window, information or physical objects
Positions	Position in space of resources and actors.
Activities	
Basic actions	A finite number $n$ of basic actions in the social practice. Joint plans are
	built using these actions.
Capabilities	A function that indicates the set of actions an actor is capable of per-
*	forming.
General preconditions	Conditions assumed to hold at the start of the practice.
Meanings	
Purpose	The purpose of an action or interaction is the reason it is performed
*	and usually indicates the intended result. E.g. a query is posed with
	the purpose to get relevant information.
Promote	Indicates that in the context of this social practice, an action $\alpha$ pro-
	motes value $v$ . E.g. waiting for your turn promotes cooperation. For a
	more formal characterization of the promotes relation see [20].
Counts-as	Indicates that within this social practice executing action $\alpha$ is seen as
	performing $\beta$ . E.g. summarizing the dialogue counts-as moving to the
	end of the dialogue
Expectations	
Plan pattern	The plan pattern indicates the landmarks (or goal state) of each part
	of the interaction, the actor/role having initiative for each part and the
	order in which the parts are executed.
Norms	There are Obligations and Prohibitions. E.g. the chatbot cannot end
	the dialogue.
Triggers	Indicates that if a condition holds then all actors believe some specific
	action will be performed next. E.g. a question to explain something is
	followed by an explanation
Start Condition	Set of conditions that actors all believe to be true when the social
	practice starts.
Duration	Expected durations of actions and plans. Note, not all actions need to
	have an expected duration.
Table 1. Description of all aspects of the social practice	

### 2.2 Argumentation

Where social practices create a context for the dialogue in which actions can be interpreted and expected, the argumentation supports the building of the actual dialogue parts. The social practice delivers so-called plan patterns that divide the whole dialogue into a number of phases each with its own goal state. These goal states are the conditions that have to be argued for using the argumentation theory.

In many open dialogue situations an argumentation dialogue with an associated argument graph is a better basis for dialogue management than decision trees coupled with traditional databases. An argument graph is a formal representation of statements that support and reject claims [6, 9]. Such a formally defined graph makes it possible to store and structure argumentation dialogues and the arguments put forward in these discussions [5], as well as the reasoning over these arguments [16, 3]. Argumentation dialogues are a method to exchange information efficiently by only requesting information that is necessary to support some main claim of the dialogue given the history of the dialogue. For instance, if the main claim is whether or not an incident is a case of fraud, then argumentation dialogue management helps us to determine what information is relevant with respect to supporting or rejecting that claim. If from the reasoning over the arguments in the graph it is clear that the status of the main claim cannot change (supported or rejected) then we may terminate the dialogue. In other words, argumentation dialogues support both reasoning with structured arguments and dialogue coherence [15]. Reasoning with arguments is very useful for further processing the results of the dialogue (cf. [4]), and is not easily achievable with decision tree based dialogue management. Argumentation dialogue systems are also easier to maintain: we can easily add or remove arguments without the need to update a complex decision tree on when to ask which information. Finally, a statement with multiple facts is relatively easy to process since we can add all facts to the graph and reason about the arguments that they support or attack.

# 3 Social Practices and Argumentation for Dialogue Management

In this section we will show how the intuitions of social practices and argumentation can be concretely used to implement a dialogue management system. The system is based on the SALVE system [1] used to train medical students to have conversations with patients. The following sections describe the modules using as example a doctor consult to diagnose a patient. The purpose of this practice is the arrangement of an health care plan for the patient. The plan assumes the accomplishment of dialogue activities such as greetings, health problem description and therapy acquisition.



Fig. 1. The agent architecture

### 3.1 The agent architecture

We assume that the agent has already identified the current social practice [1]. This is accomplished by means of the *Interaction* and *Representation and Interpretation* (R&I) modules (see figure 1). The architecture also includes a score module used to assess the utility of the dialogue. At the end of the consultation the doctor should have a clear diagnosis of the illness of the patient and a treatment plan. The patient should understand what is wrong with him and have all information needed to follow the treatment plan. If the goal is arguably achieved the dialogue gets a full score, otherwise the utility decreases. The argumentation system can be used to check the score against the argumentation tree that is built during the dialogue and arguments can be given of why a full score is not reached (yet).

Interaction Module The interaction module is composed of a graphical interface where an avatar is used to express the emotions and reactions of the system, and a *dialogue* sub module that allows the agent to manage the conversation, interpret the utterances of the players and recognize events. The dialogue sub module is implemented by means of the S-AIML technology [2], an extension of the AIML language [12], that allows to bind the categories to specific practices and their activities. The S-AIML language and its relative processor, are enhancements of the traditional dialogue engine of Alice that allows the chatbot to manage the dialogue according to a specific social practice. The language is used to describe possible question - answer modules, called categories, that can be matched during the conversation. The categories can be matched in different

```
a)
coategory>
<pattern>
<pattern>
</pattern>
</pat
```

Fig. 2. S-Aiml examples

practices (*Common Categories*), or can be organized according to the current social practice (*SP Categories*) by means of a *social practice* tag or more specifically according to a scene of the practice (*Scenes Categories*) by means of the (*scene* tag). Moreover it is possible to indicate specific preconditions that have to be matched using the (*precondition* tag). The category shown in figure 2-a allows for an interpretation of a greeting as the accomplishment of the subgoal associated to the greeting scene, and as result triggers a *GreetingReceived* event.

The category shown in figure 2-b allows for a transition to the next scene of the practice. In particular, if the user asks the agent the reason of the appointment, there is a transition to the *health problem* scene and the querying of the R&I module to get the specific health problem of the scenario. The answer of the agent will be: "*I'm experiencing one of my usual headaches. It's quite strong*".

**Representation and Interpretation** This module organizes the knowledge and the rules required by the agent to correctly interpret the events recognized by the interaction module. Similarly to the previous module, there is general knowledge and knowledge specific for the social practices. The general knowledge is related to the agent beliefs and it also includes the formalization of the social practice module, its main components, the rules for starting and closing the practices and to choose and perform a plan pattern. Taking the doctor consult as example we briefly sketch some of this knowledge according to the social practice.

From the **context** information of the social practice we get the knowledge of who the *a*ctors are and what their *r*oles are. In the example the avatar will be a patient and the student plays the doctor role. For a specific session the patient will be a specific patient with name, history, etc. and the user will be asked for her profile (name, age, etc.). We also get the resources and positions form the social practice context. In the example the patient's health file is a resource, but it can also be a form that is filled in jointly during the dialogue. Positions are not really important in this scenario, but can be helpful when more parties are involved in the dialogue.

The **activities** of the social practice indicate the activities that the participants can execute. In dialogues these are of course dialogue acts, but also can be actions like pointing or filling a form and also can indicate mental activities like making a diagnosis by the doctor. With the activities we also store their effects (both social as well as physical). The capabilities specify for each role which capabilities it has such that it can perform certain actions. Thus a doctor can make a diagnosis, while the patient is not expected to do this. In the other hand the patient can tell his symptoms, while the doctor has no direct access to them. This modularizes the type of acts that are expected from each participant in the dialogue. Finally, the general preconditions indicate some common knowledge that can be used and is assumed to be used by all parties in the interaction. E.g. the patient comes to get a consult with the doctor and has some health problem.

The **meanings** of the social practice are used for giving direction and interpreting what happens in the dialogue. The whole social practice has a purpose. In this case that the patient receives the proper health care plan. This goal is checked whenever one of the parties wants to finish the dialogue in order to see if the goal is reached. If not, the dialogue is resumed again. Also the different phases in the practice have their own purpose. E.g. the information gathering phase has as purpose for the doctor to have enough information to make a diagnosis. Thus the dialogue is focused on reaching this goal and it is explicitly checked at each point of the phase. It leads to selection of dialogue acts that contribute to this goal. The counts-as information in the social practice indicates the way actions are interpreted in the context of the practice. Thus if the doctor says "let me summarize", she does not just indicates what she is going to do, but also indicates that she thinks there will be no further discussion or questioning and the consult is going to finish. Finally, the promotes information indicates that certain ways of acting promote certain values. This will be used to make choices between reactions in order to promote a certain value or generate emotions.

The **expectations** of the social practice are more actively used in the dialogue management. The plan patterns are used to provide the contexts of the interactions. Thus they limit what is expected in each phase of the dialogue. The triggers indicate some standard behavior that does not require any processing. E.g. saying "hello" at the beginning of the dialogue triggers a greeting. These reactive rules make the dialogue processing more efficient as it can be handled by the chatbot rules directly without having to go through a more extensive reasoning phase. The norms are used to select standard actions in situations where they are expected. Thus a doctor is expected to start the consult with an inquiry about the health of the patient. Even though she might already see what is wrong if symptoms are clearly visible. Finally start condition and duration are two types of information that are used to select options at the start and end of a dialogue (phase). Generally one knows that a doctor consult will take no more than 20 minutes. Thus the expectation is that a doctor will direct the dialogue to the conclusion when this time limit is reached. Timing can thus be used to direct action selection and also influences the general state of the dialogue partners. E.g. if an answer takes a long time to come, inferences are made by the doctor or the patient about the reason for this hesitation.

In conclusion we can say that a specific social practice carries implicit knowledge, therefore, depending on the selected social practice, the agent updates its beliefs about the social interpretation of the elements with which it interacts.



Fig. 3. Drools examples

The knowledge carried by a practice includes the agents beliefs and expectations in a practice, the rules for the generation of plans and the analysis of possible norm violations. The rules embedded in a practice are used as a kind of short cuts in the deliberation of the agent. The events triggered by the conversational moves of the player, are analyzed according to the dimensions of interest in the dialogue. As a specific example, a greeting received in time is interpreted as expected relevant desirable event according to the OCC classification of emotions [14] (figure 3-a) and will elicit positive emotions (figure 3-b). In the figure one can see that information from the chatbot is exported to Drools<sup>1</sup> which then produces values for parameters that are used to generate output of the chatbot again.

We can also use Drools rules to check for information in the conversation by extracting facts from utterances using natural language processing and building an argument graph from those facts in order to determine whether, for example, enough information is provided about the health care plan. The system then asks the student through a single turn interaction to provide extra information. This way we can see how students describe certain facts. For instance, if the patient asks "Why did you decide to prescribe spironolacton?", we may gather natural language descriptions of reasons to prescribe that medicine. For example the responses may include "Because I think spironolacton is important to support your heart" or "Because the spironolacton is needed to complement the beta blocker". From these replies we may retrain our fact extraction module in order to extract information from explanations and extract also more details.

From the extracted facts we need to construct an argument graph that either supports or rejects the claim whether the patient can understand the reasons for taking all medicines. This requires a knowledge base to model argumentation in the target domain. The knowledge base about the target domain is to be acquisitioned from the experts in the domain. It contains argumentation patterns such as "spironolacton is used to support the heart functioning effectively". The argument graph is constructed by applying the argumentation patterns on the facts. The system can see what information from the student could make the difference between the main claim being supported or rejected. This missing information forms a basis for the questions that the system may ask to the student. The responses provide extensions to the argument graph. Hence, if we

<sup>&</sup>lt;sup>1</sup> Drools is a tool to describe expert system rules. See www.drools.org for documentation and software.

can translate a dialogue to an argument graph, and we specify a utility for the argument graph, then we can use this as the utility of the dialogue. The utility ought to be maximal if the main claim is supported or rejected and no possible information can be given by the student to change the status. Otherwise the utility should deteriorate depending on the information that could still be provided.

On the implementation side we can use the AIF ontology [6] to store argument graphs and the knowledge base. The graphs are evaluated (to determine what claims are supported or not by the argument graph) using Aspic<sup>+</sup> semantics [16, 3] and make use of the TOAST [19] tool.

#### 3.2 User status

In the previous sections we have explained how the conversations with the user are managed using social practices. We also showed how argumentation graphs can be used to determine whether the dialogue has been useful. However, we have not yet discussed the state of the complete dialogue from both the patient and the student perspective. In this section we will further discuss how this approach supports the "scoring" of a conversation. I.e. the overall success of the dialogue as perceived by its participants.

Our approach using social practices and mental states of the agents suggests a scoring approach that is not tied to the actual conversation moves (as is done in some dialogue games like Communicate! [11], but rather to the mental states that the agents maintain and the deviations or confirmation of the social practice that is followed. As we have shown, the user's conversational moves lead to changes in e.g. the emotional state of the conversational agent depending on the specific content of that move. Thus the target of the system is to communicate in a way to optimize the score of the user's mental state. In a naive approach one might then say that a move that makes a user angry, worried or impatient is bad (leads to a decrease in score). However, it is not always possible to avoid negative emotions in a user. If the bank finds out that the user most likely cannot get the mortgage she wants the bank needs to tell this and this might create all kinds of negative emotions in the user. The dialogue management should be geared to how to cope with these emotions and use the conversation to decrease these emotions in some way in the user. If a user is both worried and angry it can be best to first alleviate his worries and after that on his anger, but if they are tackled in the reverse order it might still work in the end. Thus scores should not be given to single states, but rather to sequences of states. E.g. a user gets a negative score if the perceived anger of the user stays above a certain level more than five or ten moves and even gets a very negative score if the user is still angry at the end of the conversation. So, rather than scoring individual states we can give rules on the sequences of states that appear and their influence on the score. A second issue is that there are several aspects that play a role in the conversation and that the system should try to control. We have mentioned empathy with the user in our context. Because the (felt) empathy is modeled as a separate aspect of the state of the user it can get a separate score. Having to handle combinations of aspects the system now has to try to maximize the scores on all aspects, which often cannot be done at the same time (i.e. in one move). By attaching the scores to rules over state sequences it is now easily possible to show that some orders of conversational moves lead to better scores than other ones without limiting the conversation to some predetermined order(s). So, we get more flexibility and more accurate scoring of the systems performance. The social practice is used for the scoring as this practice gives a context in which the interaction takes place and it therefore also indicates the expected (range of) mental states of the participants during the conversation. Thus we can use these mental states as a kind of target values for the system to achieve. Thus scores become relative to a social practice rather than some absolute values, which makes the scoring more realistic. E.g. a user should never be angry in a regular meeting with his bank account manager, but he can be expected to become angry in a bad news conversation. Thus the user becoming angry is not always bad, but is bad in certain contexts.

# 4 Discussion and Conclusion

In this paper we have discussed some new approach to dialogue management using chatbots in combination with social practices and argumentation theory. Although chatbot technology has progressed a lot and it is possible (also through learning) to generate reasonable responses to a wide range of inputs it still is inadequate for many domains. We have taken an example from the medical domain as this shows very well that responses in some domains should really fulfill many requirements. When a student is learning to manage a dialogue with patients in an anamnese consult the patient (played by the computer system) should be believable and it is very important that the student learns to take emotions of a patient into account as well as whether a patient really understands his situation, whether the patient understands the health care plan, etc. So, in this situation the patient should not just give any response to input of the student (trying to act the doctor role), but responses that indicate to the student that certain information is still missing and why it is needed as well.

The use of social practice theory provides a way to contextualize the dialogue and thus providing all kinds of background information that can be subsequently used to interpret the input of the student and also to guide the conversation. By modularizing the interaction it prevents the explosion of possible combinations of dialogue moves between the parties. By using a kind of state parameters we can also interpret dialogue moves of the student in a better way.

Argumentation is used to keep track of whether the goal of a certain dialogue phase has been achieved. The argumentation tree that is constructed can be used to prove that certain aspects have not been covered yet and it can also be used to explain why those parts are important. Moreover the tree can be used after the dialogue is finished to learn about what happened in the dialogue and to follow it up with other actions.

In this paper we have only briefly sketched how both argumentation theory and social practices together can be used to enrich chatbot technology. We have already built working systems for different applications using them separately. This paper is a first step in showing how they can be combined in a more powerful and robust dialogue management system. The next steps are twofold. First we will actually combine the different parts of the implementations into a prototype that can be used for a testbed. Secondly, we will start developing management tools to support the description of social practices and ontologies needed for applications of the system in actual real-time situations. A longer term development will be to introduce machine learning techniques to refine both the ontologies for an application domain as well as the social practices. This will lead to less development effort, but also to an adaptation to the use of the tool in specific domains and for specific types of users. Thus a dialogue can be adapted when interacting with younger children or older people, or with people with different cultural backgrounds and expectations. All these groups use different language and have different expectations on how they are addressed. This can be accommodated by the system while being based on the same basic interaction pattern.

### References

- 1. Agnese Augello, Manuel Gentile, and Frank Dignum. Social practices for social driven conversations in serious games. In *International Conference on Games and Learning Alliance*. Springer International Publishing, 2015.
- Agnese Augello, Manuel Gentile, Lucas Weideveld, and Frank Dignum. A model of a social chatbot. In *Intelligent Interactive Multimedia Systems and Services*, pages 637–647. Springer International Publishing, 2016.
- Floris Bex, Sanjay Modgil, Henry Prakken, and Chris Reed. On logical specifications of the Argument Interchange Format. *Journal of Logic and Computation*, 23(5):951–989, 2013.
- 4. Floris Bex, Joeri Peters, and Bas Testerink. A.I. for Online Criminal Complaints: From Natural Dialogues to Structured Scenarios. In A.I. for Justice workshop (ECAI 2016), 2016.
- Floris Bex and Chris Reed. Dialogue templates for automatic argument processing. In Proceedings of COMMA 2012, Frontiers in Artificial Intelligence and Applications 245, pages 366–377, 2012.
- Carlos Chesñevar, Sanjay Modgil, Iyad Rahwan, Chris Reed, Guillermo Simari, Matthew South, Gerard Vreeswijk, and Steven Willmott. Towards an Argument Interchange Format. *The Knowledge Engineering Review*, 21(04):293–316, 2006.
- 7. H.H. Clark. Using Language. Cambridge University Press, Cambridge, MA, 1996.
- 8. V. Dignum. A Model for Organizational Interaction: based on Agents, founded in Logic. SIKS Dissertation Series 2004-1. Utrecht University, 2004. PhD Thesis.
- 9. Thomas F Gordon, Henry Prakken, and Douglas Walton. The carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10):875–896, 2007.
- 10. G. Holtz. Generating social practices. JASSS, 17(1):17, 2014.
- 11. Johan Jeuring, Frans Grosfeld, Bastiaan Heeren, Michiel Hulsbergen, Richta IJntema, Vincent Jonker, Nicole Mastenbroek, Maarten van der Smagt, Frank Wijmans, Majanne Wolters, and Henk van Zeijts. *Communicate! – A Serious Game for*

Communication Skills, pages 513–517. Springer International Publishing, Cham, 2015.

- M. D. G. B. Marietto, R. V. de Aguiar, G. D. O. Barbosa, W. T. Botelho, E. Pimentel, R. D. S. Frana, and V. L. da Silva. *Artificial Intelligence Markup Language: A Brief Tutorial.* arXiv preprint arXiv:1307.3091, 2013.
- 13. Peter McBurney and Simon Parsons. Dialogue games for agent argumentation. In *Argumentation in artificial intelligence*, pages 261–280. Springer, 2009.
- Andrew Ortony, Gerald L. Clore, and Allan Collins. The Cognitive Structure of Emotions. Cambridge University Press, 1988.
- Henry Prakken. Coherence and flexibility in dialogue games for argumentation. Journal of logic and computation, 15(6):1009–1040, 2005.
- Henry Prakken. An abstract framework for argumentation with structured arguments. Argument and Computation, 1(2):93–124, 2010.
- 17. A. Reckwitz. Toward a theory of social practices. European Journal of Social Theory, 5(2):243–263, 2002.
- E. Shove, M. Pantzar, and M. Watson. The Dynamics of Social Practice. Sage, 2012.
- Mark Snaith and Chris Reed. TOAST: online aspic<sup>+</sup> implementation. In Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012, pages 509–510, 2012.
- 20. T. van der Weide. Arguing to Motivate Decisions. Utrecht University, PhD thesis, 2011.