Engagement driven Topic Selection for an Information-Giving Agent

Nadine Glas
Institut Mines-Télécom,
Télécom ParisTech, CNRS, LTCI
glas
@telecom-paristech.fr

Ken Prepin
Rakuten Institute
of Technology - Paris
ken.prepin
@priceminister.com

Catherine Pelachaud
CNRS, LTCI,
Télécom ParisTech
pelachaud@
telecom-paristech.fr

Abstract

We propose a model for conversational agents to select the topic of interaction in agent-initiated information-giving chat. By taking into account the agent's dynamically updated perception of the user's engagement, the agent's own preferences and its associations between topics, the agent tries to select the topic that maximises the agent and user's combined engagement. The model offers engagement driven dialogue management on the topic level.

1 Introduction

Conversational agents often employ a strict task-oriented dialogue structure in order to achieve the particular task for which they are built. Chat-based systems on the other hand, allow for less rigid interaction but the agent has less control of the topic of the interaction. Some applications however, ask for dialogue that falls in between these categories: Where there is not a clear task to achieve and where the interaction is not completely open either, but where there is freedom of topic choice within a certain domain. We are interested in the latter category, more specifically in interaction that is not task-driven but instead driven by social variables of the interaction. In this work, we present a topic selection model for a conversational agent, driven by the social variable *engagement*.

By taking into account the agent's perceived (detected) level of user engagement as well as the agent's own preferences and associations in the selection of a topic, we do not consider the dialogue merely from a user-oriented system point of view, but consider the agent as an interaction participant with human-like features that contributes to the interaction from its own point of view. In Section 4 we will detail the exact interpretation of these variables.

We consider engagement as *"the value that a participant in an interaction attributes to the goal of being together with the other participant(s) and of continuing the interaction"* (Poggi, 2007). In order to favour the user's engagement level previous research manipulated the agent's non-verbal behaviour including gaze (Peters, 2005), gestures (Sidner and Lee, 2003), postures (Peters, 2005) and facial displays (Bohus and Horvitz, 2009), as well as the agent's verbal behaviour including the form (Glas and Pelachaud, 2014) and prosody (Foster, 2007) of its dialogue strategies. As mentioned above, certain interaction types however, also allow for an adaptation regarding the content of the agent's dialogue strategies. In this work we focus on the latter, by proposing a model where the agent initiates discussion topics that are adapted to the user.

In the following section we will first further specify the type of interaction and topic we are looking at. In Section 3 we present related work and in Section 4 we introduce the variables that will be taken into account in the topic selection model. In Section 5 we present the topic selection model itself. In Section 6 we discuss its configurations and in 7 its implementation. Section 8 concludes our findings.

2 Information-Giving Chat

The work we describe in this paper is conducted in the context of the French project 'Avatar 1:1' that aims at developing a human-sized virtual agent playing the role of a visitor in a museum. The agent's task is to engage human users in one-to-one face-to-face interaction about the museum and some of its art objects with the objective to give the visitors information about these subjects. The choice of the exact subject is secondary: what matters is that some amount of cultural information is transferred. We refer to this type of interaction as an *information-giving chat* (as opposed to information-seeking chat (Stede and Schlangen, 2004)). Just as information-seeking chat (Stede and Schlangen, 2004), information-giving chat is distinguished by its more exploratory and less task-oriented nature, while still being more structured than general free conversation.

The information-giving chat that is modelled in this particular project is agent-initiated in order to increase the likelihood of understanding the user's contributions. Due to the limitations of our natural language understanding module it is also the agent who introduces (initiates) the topics in the interaction.

The notion of *topic* in interactions can mean different things (Brown and Yule, 1983). We define topic from a discourse perspective as *what is being talked about in a conversation* (Brown and Yule, 1983). In the context of the information-giving chat defined above, each topic refers to the discussion phase of an artwork in the museum (O'Donnell et al., 2001). Each topic is thus associated to a fragment of conversation (similar to Macias-Galinde et al. (2012)) consisting of at least 1 pair of agent-user turns. Subtopics are subfragments of these larger conversation fragments and discuss a particular aspect of the artwork. For example, the artist of the artwork or the historical period during which it was created.

3 Related Work

Some previously built virtual agent systems give their users the opportunity to directly select or reject the topics of interaction (Bickmore et al., 2011; Kopp et al., 2005), thereby adapting the content of the interaction to the user. However, these systems only offer the user a choice for certain information. They do not present a conversational virtual agent than can select interaction topics itself based on dynamic social variables in the interaction.

In order for the agent to be able to select appropriate interaction topics itself, it needs to dispose of a domain knowledge representation mapping to the possible discussion topics. Several dialogue systems dispose of some kind of representation of domain knowledge, developed for various modules such as natural language understanding (Milward and Beveridge, 2003), topic tracking (Carlson and Hunnicutt, 1996; Jokinen and Wilcock, 2012), question-answering (Agostaro et al., 2005), response generation (Pilato, 2011), surface realization (Milward and Beveridge, 2003), and the selection or generation of dialogue topics (Chakrabory et al., 2007; Macias-Galindo et al., 2012; Stede and Schlangen, 2004). We are interested in the latter where domain knowledge is organised as in such a way that it represents (potential) interaction topics.

The topic representations can be divided in specific task-oriented models (Chakrabory et al., 2007) and non task-oriented models. As information-giving chat has a less task-oriented structure (Section 2) we focus on the latter category. In this category Macias-Galindo et al. (2012) use a semantic relatedness mechanism to transition between conversational snippets in an agent that engages in chatty dialogue, and Stede and Schlangen (2004) use an ontology-like topic structure that makes the agent produce coherent topic follow-ups in information-seeking chat. However, these systems do not take into account the user's engagement during the different discussion phases (topics). They are merely oriented towards dialogue coherence and are therefore not sufficient for an optimisation of engagement by topic selection.

Song et al. (2009) and Adam et al. (2010) do take into account the user's interests or engagement level in that they decide when the agent should switch topic. The systems are in charge of the timing of a topic change. The new topics are then respectively extracted from the web or a topic structure. For the selection of the topics themselves the user's engagement or preferences are not taken into account.

By using some concepts of the models described above, we aim at building a topic structure in the agent's mind to retrieve dynamically, during human-agent information-giving chat, engaging interaction topics. Opposite to existing topic selection systems that have focused exclusively on dialogue coherence, our topics will be generated from an agent perspective: The topic structure is representing a part of the agent's knowledge, which is located within the agent's mind, and the agent's objective is to constantly favour engagement. As such the topic selection will include human-like features by taking into account the agent's dynamically updated perception of the user's engagement, the agent's preferences and the agent's associations with respect to the current topic of conversation. In the section below we define these variables.

4 Variables for Topic Selection

In order to select engaging discussion topics, the agent needs to be able to predict the *user's engagement level* during the discussion of sofar unaddressed topics (objects). For this we need to know if there are any underlying observable *preferences* that can help the agent collect indications with regard to its prediction of the user's engagement. We interpret a preference as *"a relatively stable evaluative judgement in the sense of liking or disliking a stimulus"* (Scherer, 2005). Since a topic of conversation in our interaction setting corresponds to the discussion of a particular artwork, we verified by means of a perceptive study if there exists a relation between the user's engagement level during the discussion of an artwork with a virtual agent, and the user's preference for the physical artwork that is discussed. Below we shortly describe this study (for details see Glas and Pelachaud (2015)).

4.1 User Preferences and Engagement: Perceptive Study

We simulated a small museum in our laboratory by hanging photos of existing artworks on the walls. The artworks were chosen as to vary in style and type of affect they might evoke. When the participant finished observing the artworks in a first room, the visit continued in the next room where the participant talked with Leonard, introduced as a virtual character who also visits the museum. In the interaction Leonard discussed the different artworks from the museum in a random order.

After the interaction we presented the participants a questionnaire in which we asked indirectly for the user's engagement level during the different discussion phases, corresponding to each separate discussion around a museum object. We also asked for the user's preferences of the physical artworks.

Analyses of the data collected from 33 participants (13 female, aged 19-58) regarding the randomly discussed artworks have shown amongst others that the user's preference for a museum object is significantly, positively correlated with the user's engagement (*wanting to to be together with Leonard* $p < 0.001$, $\tau = 0.50$; *wanting to continue the interaction* $p < 0.001$, $\tau = 0.52$) during the discussion of this object with a virtual agent.

From this finding we can derive that the user's

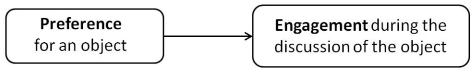

Figure 1: A human's preference and engagement.

preference for a physical object gives a direct indication of the user's engagement level during the discussion of this object (schematised in Figure 1). This makes that the characteristics (i.e. attributes) of a physical object can help the agent predict the user's future engagement level for the discussion of the object (further discussed in Section 4.3). The agent can then use its *predicted level of user engagement* for every object discussion to select an engaging topic of conversation.

4.2 Agent Preferences and Engagement

To represent human-like features in an agent that plays a museum visitor, the agent needs to have its own preferences for the artworks as well, as representing agent preferences is fundamental for any agent model (Casali et al., 2011). Besides, the preference representation of the agent can be used to express (consistent) agent appreciations, which has shown to significantly favour the user's engagement (Campano et al., 2015).

Following the correlation we found above (Section 4.1, Figure 1) an agent likes to talk most about its preferred topics as those maximise its own engagement. However, the agent we model also wants to engage the user. The agent thus tries to optimise the engagement level of both the user and the agent itself (from here onwards indicated as combined engagement). To achieve this, for each (sub)topic (object and characteristic) that can be addressed the agent calculates an *expected (predicted) level of combined engagement* and selects the one with the highest score as the next topic of discussion. In this way, the agent selects a new topic of conversation based on a combination of the agent's own preferences for the artworks and its prediction of the user's level of engagement during the discussion of the artworks. Figure 2 shows this relation.

4.3 Associations between Topics

A last human-like variable that needs to be represented when the agent selects a topic of conversation in information-giving chat are its own associations between topics. This is needed since events that share meaning or physical similarity become

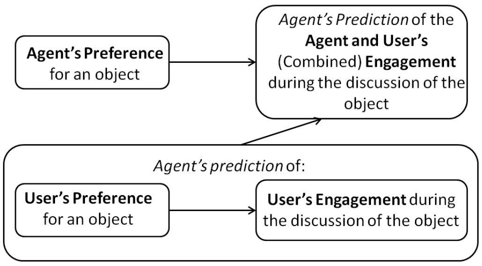

Figure 2: The agent's prediction for the level of combined engagement during the discussion of an object depends on several variables.

associated in the mind (Dellarosa, 1988). Activation of one unit activates others to which it is linked, the degree of activation depending on the strength of association (Dellarosa, 1988). The discussion of one topic can thus be associated with other topics in the agent's mind by means of similarities or shared meanings between the topics. In the context of information-giving chat about museum objects each topic is revolving around an artwork. The topics can thus be associated in the agent's mind by similarities between the physical artworks. For example, an abstract painting by Piet Mondriaan may be associated with other abstract paintings, and/or with other works by Piet Mondriaan. In the topic selection model we will therefore represent the agent's associations by similarity scores between every pair of physical objects underlying two topics.

The associations (based on object similarities) allow the agent to make predictions about the user's engagement during sofar unaddressed topics: When the user has a certain engagement level during the discussion of the current topic (and thereby a related preference towards the current object under discussion (Section 4.1)), similar conversation topics are expected to have similar levels of user preference and are thus expected to lead to similar levels of user engagement (Figure 1). The topic selection model described in the following Section ensures that when the agent's predicted user engagement level for an associated topic is high enough (in combination with the agent's own preferences) it is a potential new topic of conversation, triggered by the agent's associations.

5 Topic Selection Model

In the spirit of (Stede and Schlangen, 2004) we define an ontology-like model of domain knowledge holding the conceptual knowledge and dialogue history. The model is part of the agent's knowledge and dynamically enriched with information representing the variables described above (Section 4).

The topics all consist of artwork discussions and are therefore not hierarchically ordered but represented in a non-directed graph $\{Obj, Sim\}$ (e.g. Figure 3) where each node represents an object among N objects: $\{Obj_i, i \in [1-N]\}$.

Each object node contains the object's name (corresponding to a topic) and its characteristics (attributes) that map to the topic's subtopics (see Section 2): $\{Char_n(Obj_i), n \in [1-C], i \in [1-N]\}$, where C is the number of characteristics for any object Obj_i.

All the topics are connected to each other by similarity scores: $\{Sim(Obj_i, Obj_j), i, j \in [1-N, i \neq j]\}$ (ranging from 0 to 1), which are responsible for the possible associations of the agent. Likewise, all the subtopics (characteristics of the objects) are connected to each other: $\{Sim(Char_n(Obj_i, Obj_j), i, j \in [1-N, i \neq j]\}$.

For every object and characteristic the agent has its own preferences: $\{Pref_a\ (Obj_i), Pref_a(Char_n(Obj_i)), a = agent\}$, where 0 corresponds to no liking and 1 to a maximum liking, following the definition in Section 4. The agent also has for every object and characteristic a continuously updated predicted level of the user's engagement during the discussion of these objects and characteristics at time $t+1$: $\{Eng^*_u(t+1, Obj_i), Eng^*_u(t+1, Char_n(Obj_i)), u = user\}$, where 0 refers to the minimum level of engagement to continue an interaction and 1 refers to the maximum level of engagement.

The latter two variables lead to a continuously updated predicted level of combined (user and agent) engagement by the agent for each object and characteristic for time $t+1$: $\{Eng^*_{u+a}(t+1, Obj_i), Eng^*_{u+a}(Char_n(t+1, Obj_i)\}$, ranging from 0 to 1. See Figure 3 for a graphical representation of the topic structure that incorporates all the variables of the (sub)topics.

For $\forall Obj_i, i \in [1-N]$ the agent's predicted level of combined engagement at time $t+1$ (described in Section 4.1) during the discussion of any Obj_i is:

$$\begin{aligned} &Eng^*_{a+u}(t+1, Obj_i) = \\ &w(t) \cdot Pref_a(Obj_i) + \\ &(1-w(t)) \cdot Eng^*_u(t+1, Obj_i) \end{aligned} \tag{1}$$

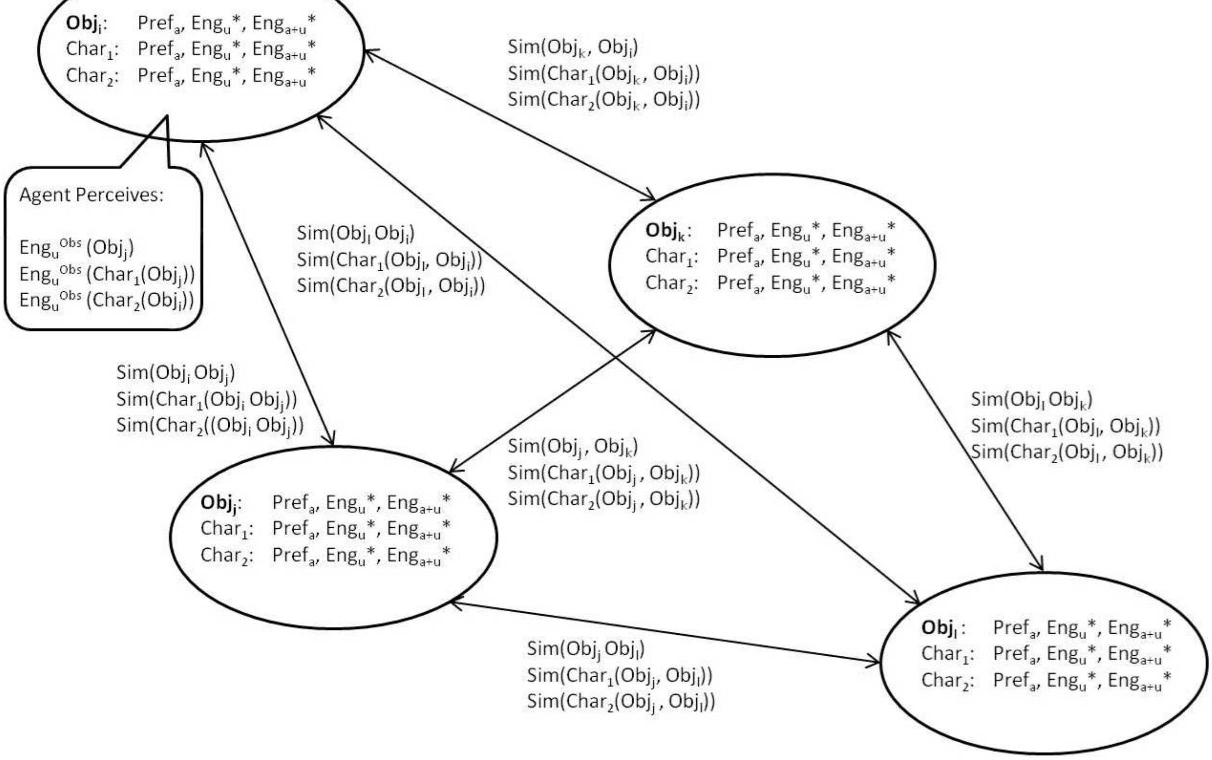

Figure 3: The graph representing an example of a topic structure in the agent's mind at a time t. Each circle represents a topic (object) where Obj_j is the current object under discussion.

Where $w(t)$ is the ratio indicating to what extent the agent values its own preferences in comparison to the user's engagement at moment t. The same equation holds for $\forall Char_n(Obj_i), n \in [1-C], i \in [1-N]$ by replacing Obj_i by $Char_n(Obj_i))$.
The agent's prediction of the user's future $(t+1)$ engagement during the discussion of Obj_i (and its characteristics by replacing Obj_i by $Char_n(Obj_i))$) is:

$$\begin{aligned} &Eng_u^*(t+1, Obj_i) = \\ &Eng_u^{obs}(t, Obj_j) \cdot Sim(Obj_j, Obj_i) + \\ &Eng_u^*(t, Obj_i) \cdot (1 - Sim(Obj_j, Obj_i)) \end{aligned} \tag{2}$$

Where $Obj_i \neq Obj_j$ and $Eng_u^{obs}(t, Obj_j)$ is the agent's observed level of user engagement during the discussion of Obj_j at time t.

5.1 Initial State

The initial state of the topic structure used for the agent's topic selection contains all the objects and characteristics that are known to the agent. For these entities dialogue fragments have been created. For $\forall Obj_i, i \in [1-N]$ and $\forall Char_n(Obj_i)$, $n \in [1-C], i \in [1-N]$ the agent's preferences and similarity scores can be initialised at any value between 0 and 1. In case we want agent associations that correspond to observable objective similarities between objects, theoretically defined similarity measures (e.g. Mazuel and Saboutret, 2008) can be used. In the latter case the similarity score between objects can be derived directly from the similarity scores of their characteristics. We further initialise for $\forall Obj_i, i \in [1-N]$:

$$\begin{aligned} &Eng_u^*(t_0 + 1, Obj_i) = Pref_a(Obj_i) \\ &\text{and } w(t_0) = 1 \end{aligned} \tag{3}$$

The same holds $\forall Char_n(Obj_i), n \in [1-C], i \in [1-N]$ (replacing Obj_i by $Char_n(Obj_i))$). This makes that for time $t_0 + 1$ the predicted user engagement of every object and characteristic equal the agent's preferences. However, this assumption is only used as a starting point for future predictions of the user's engagement that will be based on observed user behaviour (Equation 2). At the start of the interaction (t_0), the agent only takes into account its own preferences, indicated by $w(t_0) = 1$. The first topic the agent introduces in the interaction is the one for which it predicts the highest level of combined engagement at time $t_0 + 1$:

$$max\{Eng^*_{a+u}(t_0 + 1, Obj_{1-n})\} \quad (4)$$

The agent introduces one by one the subtopics of this first topic for which:

$$\{Eng^*_{a+u}(t_0 + 1, Char_n(Obj_i))\} > e \quad (5)$$

Where e is a threshold for the minimum level of predicted mutual engagement level that the agent finds acceptable for the interaction. For example the agent can decide to only talk about the subtopics that are predicted to lead to half the maximum level of engagement, setting e to 0.5.

5.2 Updating

A new topic is selected when either: 1) The current topic is finished, meaning that the conversational fragment has been uttered completely. Or 2) the detected level of user engagement during an interval I within the discussion phase of an object is below a threshold z. The description of the user engagement detection method itself lies outside the scope of this paper. The required length of I and level of z, which determine when the user's engagement level (detected by the agent) should lead to a topic switch, will be studied in future work.

At any time t, just before selecting a new topic of interaction the agent first updates the weights in the topic structure with information that is gathered during the previous discussion phases. In the rest of this section we describe how.

For $\forall Obj_i, i \in [1 - N]$ that are part of the dialogue history (already discussed topics) we set:

$$Eng^*_u(t + 1, (Obj_i)) = 0 \text{ and } w(t) = 0 \quad (6)$$

Similarly for $\forall Char_n(Obj_i), n \in [1 - C], i \in [1 - N]$ that are part of the dialogue history. This implies that the agent makes the assumption that once a topic has been addressed the user does not want to address it again. The agent values this over its own preferences ($w(t) = 0$). This simplification makes that the system shall not discuss a topic twice.

$\forall Obj_i, i \in [1 - N]$ and $\forall Char_n(Obj_i), n \in [1 - C], i \in [1 - N]$ that are not in the dialogue history the agent's prediction for the user's engagement at time $t + 1$: $Eng^*_u(t + 1, Obj_i)$, as well as the agent's prediction for the agent and user's combined engagement level at time $t + 1$: $Eng^*_{a+u}(t + 1, Obj_i)$ are updated by Equation 1 and Equation 2. This is done by entering the agent's detected (observed) overall level of user engagement during the discussion phases of the lastly discussed object and each of its characteristics (at time t): $Eng^{obs}_u(t, Obj_j)$ and $Eng^{obs}_u(Char_k(t, Obj_j))$. This update makes sure that the agent's observed user engagement level influences the predicted (user and combined) engagement levels for the objects and characteristics that the agent associates with the previously discussed (sub)topics. As mentioned before, the detection method of the user's engagement level lies outside the scope of this paper.

In circumstances where a detection of the user's engagement is not possible Equation 1 and Equation 2 can be updated by entering the user's explicitly uttered preferences for the lastly discussed object and characteristics at the place of respectively $Eng^{obs}_u(t, Obj_j)$ and $Eng^{obs}_u(Char_k(t, Obj_j))$. This follows from the finding that a user's preference is directly related to the user's engagement (Section 4.1).

5.3 Topic Selection

Whenever a new topic needs to be introduced (see previous section) it is selected in the same way as the first topic of the interaction:

$$max\{Eng^*_{a+u}(t + 1, Obj_{1-n})\} \quad (7)$$

In this way, the agent tries to optimise the combined engagement. The selected subtopics of this topic are, like the first subtopics of the interaction, those where:

$$\{Eng^*_{a+u}(t + 1, Char_n(Obj_o))\} > e \quad (8)$$

5.4 Example

For the sake of clarity, in this Section we demonstrate the working of the topic selection model with a small example topic structure shown in Figure 4. In this example the agent knows about the 4 topics (objects) that are listed in Table 1.

Figure 4 shows how the variables for each topic can evolve over time t_{0-2} during an interaction. Due to space limitations we limit this example to the calculation and selection of topics. The calculation of the weights of the subtopics occurs in exactly the same manner. Only the ultimate selection of subtopics differs slightly as described in Section 5.3.

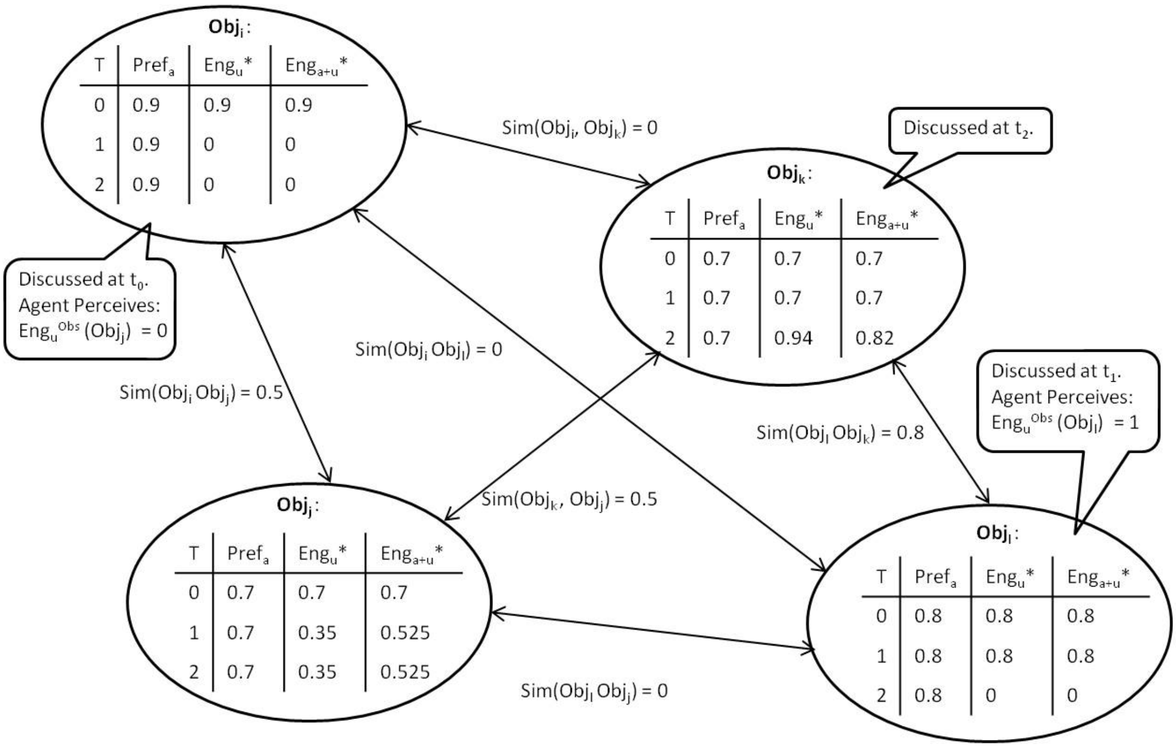

Figure 4: Example of the evolution of the weights in the topic structure over time t_{0-2}. In this example, at each t, $w = 0.5$.

Object	**Type**	**Artist**
Obj_i	Statue	Antiquity
Obj_j	Statue	17th century
Obj_k	Painting	17th century
Obj_l	Painting	18th century

Table 1: The objects of the example topic structure of Figure 4.

The values of the variables for each topic at time t_0 represent the initial state. Given that at this moment $Eng*_{a+u}$ is the highest for Obj_j, this topic is the first to be selected for discussion. When the agent then perceives a minimum level of user engagement during the discussion of this topic, the updated variables (t_1) lead to the selection of object Obj_l as next object, which has nothing in common with the former object. To show the opposite extreme situation, during the discussion of object Obj_l the agent perceives a maximum level of user engagement, leading to the selection of Obj_k as next topic, staying close to the characteristics of the former object. Of course Figure 4 is only a limited example and not sufficient to illustrate the full potential of the tradeoff between agent and user oriented variables in the selection of a topic.

6 Topic Selection Configurations

As described in Section 5.2 the preference, engagement and similarity weights in the topic structure can be initialized at any value ranging from 0 to 1. The freedom of initialising these variables as desired allows for different configurations. The initialisation of the agent's preferences, for instance, can reflect different types of agents (as recommended by Amgoud and Parsons (2002)) but can also be initialised, for example, at values that are close to the users' preferences in previous interactions. The agent's preferences for the objects can be directly related to the sum of its preferences for the characteristics of the object or not. It is also possible to attribute more importance to the preference for one characteristic than to another. The same holds for the similarity values. For example, to model an agent that is particularly focused on history, the similarity and preference weights of the characteristic "period" may have a larger impact on the similarities and preference of the entire object than the other characteristics of the object. The initialisation of the graph can be simplified with the help of a museum catalogue that already lists the objects and their characteristics.

The topic selection model can be easily extended to other domains that can be structured similarly as museum objects. This means that the agent needs to have its preferences for the different topics and can associate the topics to each other by means of similarity scores. Selecting subtopics as described in Section 5 is only possible if the topics' characteristics (attributes) can be defined.

7 Implementation and Dialogue Management

For the management of the multimodal behaviour of the agent we use the hierarchical task network Disco for Games (Rich, 2012) that calls prescripted FML files, which are files that specify the communicative intent of an agent's behaviour and include the agent's speech (Heylen et al., 2008). As Disco is developed for task-oriented interactions it offers a fixed, scripted order of task execution where agent contributions and user responses alternate. As mentioned in Section 2, in our project each (sub)topic is associated to a scripted fragment of conversation, consisting of 1 or multiple pairs of agent-user turns. Within such a conversation fragment we can thus directly use the Disco structure. The agent executes the tasks that consist of talking about the object while the local responses of the user drive the dialogue further in the network.

In between (sub)topics it is different. In information-giving-chat we cannot foresee, and thus predefine, the topics and their order of discussion as they are selected by the agent during the interaction. Therefore, at any time a topic switch is required, we overwrite the fixed task structure proposed by Disco by calling from an external module the appropriate tasks that map to the selected (sub)topics. The external module is the topic selection module of the agent. In this way we continuously paste in real time dialogue parts to the ongoing conversation.

This procedure makes that local dialogue management is controlled by Disco and topic management is controlled by the external topic selection module, thereby adding flexibility to the existing task-oriented system, resulting in a more adaptive and dynamic dialogue. The agent's topic selection module could be connected to any other task and/or dialogue system as well.

8 Conclusion and Future Work

In this work we have proposed an engagement driven topic selection model for an information-giving agent. The model avoids the need for any pre-entered information of the user. Instead, it dynamically adapts the interaction by taking into account the agent's dynamically updated perception of the user's level of engagement, the agent's own preferences and its associations. In this way, the interaction can be adapted to any user. The model's configurations also allow for different types of agents. By connecting the topic selection model to existing task-oriented systems we proposed a way to construct a dynamic interaction where the agent continuously pastes in real time dialogue parts ((sub)topics) to the ongoing conversation.

In the future we would like to perform a perceptive study to evaluate the topic selection model in human-agent interaction. However, for this we first need to plan some additional research.

First, we will study the different ways of switching topic on a dialogue generation level. Even when two consecutive topics have no characteristics in common, the agent needs to present the new topic in a natural way in the conversation without loosing the dialogue coherence (Levinson, 1983). Strategies that we will consider to achieve this include transition utterances by the agent that make the agent's associations explicit, transition utterances that recommend interesting (e.g. similar/opposite) artworks to the user, and transition utterances that refer to the artworks' locations within the museum.

Further, before evaluating the topic selection module in interaction with users, we will also have a closer look at the timing of topic transitions. We noted that a topic switch is needed, amongst others, when the current topic leads to a very low user engagement during a certain time interval. We will need to determine the exact interval of the engagement detection that is required and thereby determine the timing of a possible topic transition.

Possible extensions to the topic selection model we would like to explore in the future include the option of coming back to previously addressed topics, considering agent preferences that may change during the interaction and dealing with similar information-giving conversational fragments for different topics.

Acknowledgments

We would like to thank Charles Rich for his help with Disco. This research is partially funded by the French national project Avatar 1:1, ANR MOCA, and the Labex SMART (ANR-11-LABX-65) supported by French state funds managed by the ANR, within the Investissements d'Avenir program under reference ANR-11-IDEX-0004-02.

References

Carole Adam, Lawrence Cavedon, and Lin Padgham. 2010. Flexible conversation management in an engaging virtual character. *International Workshop on Interacting with ECAs as Virtual Characters.*

Francesco Agostaro, Agnese Augello, Giovanni Pilato, Giorgio Vassallo, and Salvatore Gaglio. 2005. A conversational agent based on a conceptual interpretation of a data driven semantic space. *Advances in Artificial Intelligence*, 381–392.

Leila Amgoud and Simon Parsons 2002. Agent dialogues with conflicting preferences. *Intelligent Agents VIII*, 190–205.

Timothy Bickmore and Julie Cassell. 1999. Small talk and conversational storytelling in embodied conversational interface agents. *AAAI fall symposium on narrative intelligence*, 87-92.

Timothy Bickmore and Julie Cassell. 1999. Small talk and conversational storytelling in embodied conversational interface agents. *AAAI fall symposium on narrative intelligence*, 87–92.

Timothy Bickmore, Daniel Schulman and Langxuan Yin. 2010. Maintaining engagement in long-term interventions with relational agents. *Applied Artificial Intelligence*, 24(6), 648–666.

Timothy Bickmore, Laura Pfeifer and Daniel Schulman. 2011. Relational agents improve engagement and learning in science museum visitors. *Intelligent Virtual Agents*, 55–67.

Cynthia Breazeal and Brian Scassellati. 1999. How to build robots that make friends and influence people. *Proceedings of the International Converence on Intelligent Robots and Systems 2*, 858–863.

Gillian Brown and George Yule. 1983. Discourse Analysis. *Cambridge University Press.*

Dan Bohus and Eric Horvitz. 2009. Models for multiparty engagement in open-world dialog. *Proceedings of the 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 225–234.

Sabrina Campano, Chlo Clavel, and Catherine Pelachaud. 2015. I like this painting too: When an ECA Shares Appreciations to Engage Users. *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 1649–1650.

Ana Casali, Llus Godo and Carles Sierra. 2011. A graded BDI agent model to represent and reason about preferences. *Artificial Intelligence*, 175(7), 1468-1478.

Rolf Carlson and Sheri Hunnicutt. 1996. Generic and domain-specific aspects of the Waxholm NLP and dialog modules. *Proceedings of the 4th International Conference on spoken Language Processing*, 677–680.

Sunandan Chakraborty, Tamali Bhattacharya, Plaban K. Bhowmick, Anupam Basu, and Sudeshna Sarkar. 2007. Shikshak: An intelligent tutoring system authoring tool for rural education. *International Conference on Information and Communication Technologies and Development*, 1–10.

Denise Dellarosa. 1988. The psychological appeal of connectionism. *Behavioral and Brain Sciences*, 11(01), 28–29.

Mary Ellen Foster. 2007. Enhancing human-computer interaction with embodied conversational agents. *In Universal Access in Human-Computer Interaction. Ambient Interaction*, 828–837.

Nadine Glas and Catherine Pelachaud. 2014. Politeness versus Perceived Engagement: an Experimental Study. *Proceedings of the 11th International Workshop on Natural Language Processing and Cognitive Science*, 135–147.

Nadine Glas and Catherine Pelachaud. 2015. User Engagement and Preferences in Information-Given Chat with Virtual Agents. *Workshop on Engagement in Social Intelligent Virtual Agents*, Forthcoming.

Fritz Heider. 1958. The Psychology of Interpersonal Relations. *Psychology Press.*

Dirk Heylen, Stefan Kopp, Stacy C. Marsella, Catherine Pelachaud, and Hannes Vilhjlmsson. 2008. The next step towards a function markup language. *Intelligent Virtual Agents*, 270–280.

Kristiina Jokinen and Graham Wilcock. 2012. Constructive Interaction for Talking about Interesting Topics. *Proceedings of Eighth International Conference on Language Resources and Evaluation*, 404-410.

Souhila Kaci. 2011. Working with Preferences. Less Is More. *Springer.*

Stefan Kopp, Lars Gesellensetter, Nicole C. Krmer and Ipke Wachsmuth. 2005. A conversational agent as museum guide-design and evaluation of a real-world application. *Intelligent Virtual Agents*, 329–343.

Stephen C. Levinson. 1983. Pragmatics. *Cambridge textbooks in linguistics.*

Daniel Macias-Galindo, Wilson Wong, John Thangarajah and Lawrence Cavedon. 2012. Coherent topic transition in a conversational agent. *Proceedings of the 13th Annual Conference of the International Speech Communication Association*, 1–4.

Laurent Mazuel and Nicolas Sabouret. 2008. Semantic relatedness measure using object properties in an ontology. *Springer Berlin Heidelberg*, 681–694.

David Milward and Martin Beveridge. 2003. Ontology-based dialogue systems. *Proceedings of the 3rd Workshop on Knowledge and reasoning in practical dialogue systems*, 9–18.

Mick O'Donnel, Chris Mellish, Jon Oberlander and Alistair Knott. 2001. ILEX: an architecture for a dynamic hypertext generation system. *Natural Language Engineering*, 7(3) : 225–250.

Christopher Peters. 2005. Direction of attention perception for conversation initiation in virtual environments. *Proceedings of the 5th International Conference on Intelligent Virtual Agents*, 215–228.

Christopher Peters, Catherine Pelachaud, Elisabetta Bevacqua, Maurizio Mancini, and Isabella Poggi. 2005. Engagement capabilities for ECAs. *AAMAS'05 workshop Creating Bonds with ECAs*.

Giovanni Pilato, Agnese Augello, and Salvatore Gaglio. 2011. A modular architecture for adaptive ChatBots. *Fifth International Conference on Semantic Computing*, 177–180.

Isabella Poggi. 2007. Mind, hands, face and body: a goal and belief view of multimodal communication. *Weidler*.

Charles Rich and Candace L. Sidner. 2012. Using collaborative discourse theory to partially automate dialogue tree authoring. *Intelligent Virtual Agents*, 327-340.

Klaus R. Scherer. 2005. What are emotions? And how can they be measured? *Social science information*, 44(4), 695–729.

Candace L. Sidner and Christopher Lee. 2003. An architecture for engagement in collaborative conversations between a robot and humans. *Mitsubishi Electric Research Labs TR2003-13*.

Xin Song, Kazuki Maeda, Hiroyuki Kunimasa, Hiroyuki Toyota, and Dongli Han. 2009. Topic control in a free conversation system. *International Conference on Natural Language Processing and Knowledge Engineering*, 1–6.

Manfred Stede and David Schlangen. 2004. Information-seeking chat: Dialogues driven by topic-structure. *Proceedings of the 8th Workshop on Semantics and Pragmatics of Dialogue*, 117–124.