

Interpreting social cues to generate credible affective reactions of virtual job interviewers

Hazael Jones, Nicolas Sabouret, Ionut Damian, Tobias Baur, Elisabeth André, Kaska Porayska-Pomsta, Paola Rizzo

Angaben zur Veröffentlichung / Publication details:

Jones, Hazael, Nicolas Sabouret, Ionut Damian, Tobias Baur, Elisabeth André, Kaska Porayska-Pomsta, and Paola Rizzo. 2014. "Interpreting social cues to generate credible affective reactions of virtual job interviewers." arXiv. <https://arxiv.org/pdf/1402.5039v2>.

Nutzungsbedingungen / Terms of use:

licgercopyright

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

Deutsches Urheberrecht

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

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



Interpreting social cues to generate credible affective reactions of virtual job interviewers

Hazaël Jones¹, Nicolas Sabouret², Ionut Damian³, Tobias Baur³, Elisabeth André³,
Kaśka Porayska-Pomsta⁴ and Paola Rizzo⁴

¹Laboratoire d'Informatique de Paris 6, 4 place Jussieu, 75005 Paris, France, hazael.jones@lip6.fr

²Laboratoire d'Informatique pour la Mécanique et les Sciences de l'Ingénieur, 91403 Orsay, France, nicolas.sabouret@limsi.fr

³Human Centered Multimedia, Augsburg University, 86159 Augsburg, Germany, {damian, baur, andre}@hcm-lab.de

⁴London Knowledge Lab, Institute of Education, London, UK, {K.Porayska-Pomsta, P.Rizzo}@ioe.ac.uk

ABSTRACT

In this paper we describe a mechanism of generating credible affective reactions in a virtual recruiter during an interaction with a user. This is done using communicative performance computation based on the behaviours of the user as detected by a recognition module. The proposed software pipeline is part of the TARDIS system which aims to aid young job seekers in acquiring job interview related social skills. In this context, our system enables the virtual recruiter to realistically adapt and react to the user in real-time.

Categories and Subject Descriptors

I.6.5 [Computing Methodologies]: Simulation and Modelling—*Model Development*; J.4 [Computer Applications]: Social and Behavioural Sciences

General Terms

Theory

Keywords

Perceptions, Social cues, Communicative performance, Affective model, Job interview.

1. INTRODUCTION

The number of young people who are not in employment, education or training (NEETs) is increasing across Europe. According to Eurostat ¹, in March 2012, 5.5 million of 16 to 25 years old European youngsters were unemployed, amounting to 22.6% of the youngster global population, which is

¹ec.europa.eu/eurostat

10 points above the entire world's population. These statistics highlight European youth unemployment as a significant problem.

NEETs often lack self-confidence and the essential social skills needed to seek and secure employment [10]. The young unemployed often find it difficult to present themselves in a good light to prospective employers, which may put them at further risk of marginalisation. Social coaching workshops, organized by youth inclusion associations across Europe, constitute a common approach to helping people in acquiring and improving their social competencies, especially in the context of job interviews. However, this is an expensive and time-consuming approach that relies on the availability of trained practitioners as well as a willingness of the young people to discuss their strengths and weaknesses in front of practitioners and often also in front of their peers.

Digital games offer a promising way of supporting the training and coaching of young people, providing them with a safe and private environment in which they can practice their skills repeatedly. TARDIS² [1] is a project funded by the FP7, whose aim is to build a scenario-based serious-game simulation platform that supports social training and coaching in the context of job interviews. It relies on the use of virtual agents that are capable of reacting in real-time to a human interlocutor's affect, based on the interlocutor's non-verbal cues that are automatically detected and analysed in the context of simulated job interviews.

The goal of this paper is to present the TARDIS' mechanisms involved in the detection and interpretation of social cues and in informing the behaviours of virtual recruiters. We show how the system builds affective states and beliefs about a real user during an interaction with a virtual character. Specifically, the TARDIS game relies on real-time social cue recognition, communicative performance computation and affective computation/decision making by the virtual recruiter. Building a credible job interview simulation involving a reactive virtual agent requires real-time information about the youngster in order to allow for an assessment of the appropriateness of the youngster's reactions

²<http://www.tardis-project.eu>

and their communicative performance. Such assessment is done through the perception of a number of relevant social cues that are interpreted in term of performance, relative to the social expectations that are associated with a given situation. This assessment allows the TARDIS system to compute both an affective reaction for the virtual recruiter and the future steps in the interaction dialogue.

The next section presents a brief overview of related work. We then present the different components of our architecture that have been implemented in the TARDIS system.

2. STATE OF THE ART

Several research projects have already considered using virtual agents to help humans improve their social skills and, more generally, their emotional intelligence [16, 34, 3]. However, while the existing approaches use a reactive approach to the user inputs, our efforts focus on deriving an affective reaction of the virtual recruiter, based on the analysis of the user's beliefs and social cues.

2.1 Social Cue Recognition

Using signal processing techniques to detect behavioural patterns is not a new idea (for an overview see [36]). However, most research to date focused on a reduced number of modalities, such as speech [37] or facial expressions [40], to infer user states, with little attention having been paid to gestures or postures [20, 22]. Furthermore, signal processing work is aimed for offline analysis of recordings, rather than real-time interactive applications. One interesting example is the public training system presented by Batrinca et al. [5], where the user is able to practice public speaking with the help of a virtual crowd. However, the behaviour analysis that they propose does not happen in real-time during the interaction, but rather post-hoc and offline. This means that the system is not able to react to the user's behaviour as they interact with the system.

There are, of course, exceptions. One recent example is the job interview simulation system presented by Hoque et al. [18]. Their system is able to detect a limited number of social cues in real-time, including smiles and audio features as well as perform speech recognition. While, owing to the large number of languages that TARDIS aims to support, our system does not implement speech recognition, it does recognise a much broader range of social cues, including bodily social cues, such as movement energy, gestures and postures, physiological features and eye gaze.

2.2 Job interview assessment based on social cues

Several researchers show that diverse interaction scenarios, including job interviews, negotiations or group meetings, are heavily impacted by the interactants' nonverbal behaviours. For instance [35] and [2] found that the assessment of candidates by job interviewers is significantly influenced by social cues such as tone of voice, eye gaze contact and body movement. For these reasons, the contemporary research in this field increasingly focuses on the relationship between nonverbal behaviours and the outcome of interactions. Specifically to job interviews, [11] studied how the success of simulated job interviews can also be predicted by conversational

engagement, vocal mirroring, speech activity, and prosodic emphasis.

The training of social skills has informed several computer-based simulation environments for domains that include bullying at school [3], intercultural communication [21] and job interviews [30]. However, in the case of job interview situations a major difficulty is that the interviewees tend to suppress their emotions and avoid showing social cues, which may result in an undesirable impression being given to the interviewer and ultimately it may impact negatively on the interview outcome. Emotional expression in job interviews is heavily governed by display rules, and most job applicants feel obliged to hide their feelings, especially if those feelings are negative [32]. The fact that in job interview situations, young job seekers tend to avoid showing emotions has been confirmed by our own studies, which will be reported in future publications. In these studies, subjects involved in both mock interviews with human recruiters and in simulated job interviews with virtual recruiters displayed very few or no bodily social cues that could be detected reliably and consistently by our current technologies. For this reason, we are currently focusing on voice cues.

2.3 Affects and Theory of Mind

Several approaches to building credible virtual humans have been proposed in the domain of affective computing³, including cognitive models of emotions [28, 25], models of personality [29] and of social relations [27]. However, to our best knowledge, no computational model of social attitudes has been proposed. Social attitudes are socially conditioned expressions of the personality of an agent as manifested in its behaviours and emotional displays. The social conditioning derives from the specific social and cultural norms of a given community or situational contexts. For example, in the context of a job interview, social attitudes provide the recruiter with information about the interviewee's personality and feelings towards the job. This information influences the way in which the virtual recruiter will progress the interview and how it will perceive the suitability of the interviewee for the given job. Many of the processes involved in interpreting the beliefs and emotions of the TARDIS' users are cognate with and draw from the studies related to Theory of Mind [26] and reverse appraisal [17].

Theory of mind (ToM) [24, 4], is the ability of a person to attribute mental states (beliefs, intentions, desires and affects) to others. Numerous studies have been conducted in relation to the reasoning process of an agent about the cognitive process of another agent [9, 13]. In our work, we want to model the reasoning process of an agent that deduces the preferences of a human (the interviewee). This particular configuration raises additional difficulties and leads to a novel formulation of the ToM. As a way of illustrating the novelty of our approach, consider for example Pynadath's approach [31], where an agent has subjective beliefs about others. Artificial agent A has beliefs about another artificial agent B, which beliefs follow the real structure of agent B's beliefs. However, since in TARDIS, agent B is human (the interviewee), we do not have access to prior access to its belief structure, which must be inferred from the outputs of

³See the Humaine project: emotion-research.net

the affect recognition module.

3. ARCHITECTURE OVERVIEW

The TARDIS architecture is composed of five main components as shown in Figure 1.

The Social Cue Recognition (SCR). This module detects and recognises various social cues that are produced by the youngsters during the simulated interviews. The module sends the detected cues to the on-line user model. To facilitate detection, the module requires the use of different sensors, such as microphones or cameras. In this paper, we will focus on the audio cues perceived with the help of a head-mounted microphone.

The Online User Model. This module receives information about the youngster's social cues detected by the SCR module, it computes a performance index based on the average values of audio social cues that can be considered suitable to a job interview, and it sends this information to the Affective Core, described in Section 6.

The Interview Scenario. This component provides the virtual recruiter with the expectations that it should have in relation to the interviewed youngster's emotions and attitudes, based on the specific stage in the interview. In the current version of TARDIS, the agent has no understanding of the youngster's actual answers to the questions. It simply follows a predetermined scenario, while focusing on the affective recognition and adaptation, based on its knowledge of the scenario and the specific interview questions that have been asked at any given point. It is well known in the virtual character research community that the definition of context-sensitive interactive scenario often requires expert programmers and typically results in a hard-to-maintain code. To avoid this, a visual authoring approach for virtual character applications is used. This tool is based on an existing authoring tool, the Visual SceneMaker that allows the scenario to be built by non computer programmers. The scenario provides a context to the virtual recruiter model, specifically – the expectations of the recruiter with respect to the desirable emotions and attitudes of the youngsters.

The Affective and Decision Model of the Virtual Recruiter. The affective and decision model represents the beliefs of the virtual recruiter about the youngster's mental states (specified in terms of affect). It also represents the recruiters intentions with respect to its possible best next interview action. The affective core periodically computes the new affective states for the Virtual Recruiter, based on the perceptions and scenario-based expectations, and current affective states of the youngsters. The difference between the expected and the actually expressed affective states of the youngster are key in the update of the virtual recruiter's own affective states (see next section). Based on the system's theory of mind of the youngster, the decision module selects the next actions in the scenario.

The Virtual Recruiter Animation. The animation module is responsible for rendering the virtual recruiter's affective state through its behaviours. To achieve realistic behaviour displays, we developed a new motion capture pipeline able to process the data captured more efficiently. The expression of affects relies on a model of interpersonal stance, which in turn is based on different time windows. This new pipeline is integrated with Greta⁴ and it uses MPEG-4, BML and FML standards to allow the user to define the agent's behaviours in XML like form. This allowed us the easy integration of the Charamel⁵ agent Gloria.

In order to determine the social cues and affects that are relevant during job interviews, a knowledge elicitation study was conducted with ten users in France⁶. These experiments involved enactments of job interviews as they are normally conducted in the participating organisation. The studies allowed us to determine the affective states of the youngsters that are relevant to this domain of interaction and that are integrated in our youngster and recruiter models. Based on the knowledge elicited from these studies, the TARDIS' dialogue model has also been designed.

The rest of this paper focuses on three important modules in the TARDIS simulation. In Section 4, we present the social cue recognition module, Section 5 focuses on the communicative performance computation, while the recruiter affective and decisional model is discussed in Section 6.

4. SOCIAL CUE RECOGNITION

The Social Cue Recognition module is based on the Social Signal Interpretation framework [38]. SSI provides an interface to diverse sensing devices, as well as a variety of tools for real-time recording and pre-processing of data of human behaviours. The functionality of the module regarding the recognition of body social cues, such as gestures, postures or facial expressions, has been presented in our previous work [12, 6, 7]. In this paper, we will focus on the detection, analysis and interpretation of the audio cues. Table 1 summarizes the relevant cues.

As a first step in the recognition of social cues, the raw data coming from the sensors is filtered and transformed with the help of an SSI pipeline and various third party libraries, such as PRAAT [8] and OpenSMILE [14]. The data streams are then forwarded to the individual social cue recognisers which trigger social cue occurrence events.

The SCR module is able to detect two types of social cues: discrete and continuous. In the case of discrete social cues, the recogniser triggers an event as soon a sufficiently strong occurrence of the respective social cue has been detected. Assessing the quality of a given social cue is either performed using a threshold approach or machine learning algorithms. The event contains information on the recognised social cue and the time of the occurrence, which is meant to be processed by the online user model that makes inferences about

⁴GRETA is an Embodied Conversational Agent used in the SEMAINE project.

⁵<http://charamel.de/>

⁶The study was conducted at Mission Locale, which is a French network of organisation that helps youngster in their social course and professional career.

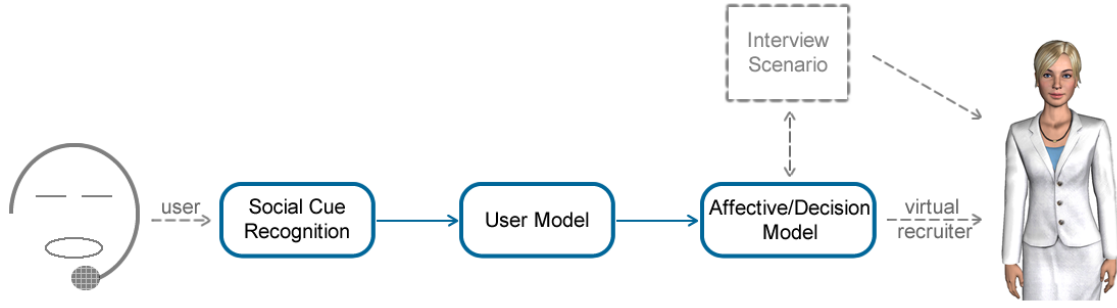


Figure 1: Global architecture

Social Cue	Description
voice activity	Presence or absence of voice
intensity, loudness, energy	Energy-based features of the audio signal
pitch value	The pitch (F0) of the audio signal
jitter, shimmer, voice breaks, harmonicity	Quality-of-voice features [8] computed from pitch
speech rate	Rate of user's speech [8]
length of speech segments	The duration in seconds of the user's speech segments determined by voice activity detection

Table 1: Audio cues recognised by the SCR module

the user's complex mental states.

Continuous social cues are sent once per second and contain information on the current magnitude or probability of the observed social signal.

5. USING AUDIO CUES TO COMPUTE COMMUNICATIVE PERFORMANCE

The online user model, currently being implemented, compares the audio cues data produced by users with a set of expected or desirable cues, i.e. audio features that are deemed suitable for each part of a job interview. Based on this comparison, the model computes an index between the values of -1 (negative performance) and $+1$ (positive performance), and sends it to the Affective Core. 0 represents neutral performance. The index is a weighted sum of the differences between each desired cue value and the actual value produced by the youngster regarding speech duration, speech rate, speech volume, pause before speech, etc.

The following are examples of expected cues:

- interviewer's questions that are complex or sensitive in nature, e.g. "do you have any weaknesses?" requires an answer that goes beyond a simple "yes" or a "no". An appropriate answer should be elaborate and thus, a short speech duration in response to a complex question such as exemplified, would be assigned a perfor-

mance index close to -1 ;

- the voice should not be too loud (intense), as it might sound aggressive, nor too quiet, as it might give the impression of lacking self-confidence;
- the speech should not be too fast, because it might be considered a sign of anxiety, nor too slow, as it might indicate boredom or disengagement;
- the interviewee should not begin to answer too long after the interviewer's questions, because this could be deemed a sign of hesitation, but equally, it the interviewee should not start answering before the interviewer finishes their question;
- the voice should have a good variability of pitch, so as not to sound boring.

Our studies of human-to-human mock job interviews have shown that there are strong individual differences among youngsters' social cues. For this reason the on-line user model requires a calibration phase to take place during an initial part of the interview (over at least the first 3 turns), and then it continuously compares the youngster's behaviours to their typical baseline as established during calibration, it identifies peaks (i.e. behaviours significantly above or below the baseline), and compares them with the expected cues in order to compute the communicative performance index.

6. AFFECTIVE & DECISIONAL CORE

The perception module and the scenario module provide the input to this module, which is a performance index of the youngster in the range of $[-1, +1]$. This index is used to compute recruiter affects and decision.

6.1 Affective Core

6.1.1 Overview of the model

The Affective Module is based on a set of rules that compute categories of emotions, moods and attitudes for the virtual recruiter, based on the contextual information given by the scenario and the detected affects (emotions, moods and attitudes) of the participant. The computation of the virtual agent's emotions is based on the OCC model [28] and the computation of the agent's moods is based on the ALMA model [15]. The details of the computation of emotions and

moods will not be presented in this paper; it can be found in [19].

The Affective Core receives a performance index as its input. The performance index, which fall in the range of $[-1, +1]$, represents the overall performance of the youngster (its attitude, its affects, its vocal performance). The detected (d) performance is denoted as P_d . Similarly, a set of expected (e) performance index is received from the scenario module. This expected performance index is linked to the difficulty of the question. For example, if a question is easy, e.g. "Did you find us easily?", the performance index will be near of 1.

Formally, in our model, all affects of the recruiter correspond to a value in the interval of $[0, 1]$ and we use \mathcal{A} to denote the set of all affects. The different affects are categorised in terms of three subsets: $E(t)$ (emotions), $M(t)$ (moods) and $A(t)$ (attitudes) are virtual recruiter's simulated affects. These emotions are computed using expert rules based on the values of $\mathcal{P}_d(t)$ and $\mathcal{P}_e(t)$. All these rules are described in [19].

The list of virtual recruiter's possible affects (emotions, moods and attitudes) that are represented in the model is given in Table 2. Note that this set is different from the affects that are actually detected and expected. It is based on the literature and on the mock interview corpus analysis (especially the knowledge elicitation phases mentioned in sections 4 and 5). The emotions are a simple subset of the OCC model [28] that was selected based on what practitioners, acting as recruiters, expressed during the mock interviews analysed. The moods originated from the ALMA model [15] are defined on 3 dimensions (Pleasure, Arousal and Dominance), but we limited them to the positive dominance zone (since recruiters do not show submissive moods in the context of job interviews). Moods of Table 2 are with positive or neutral dominance.

	Positive	Negative
Emotions	Joy	Distress
	Relief	Disappointment
	Admiration	Anger
	Hope	Fear
Moods	Relaxed	Hostile
	Exuberant	Bored
		Disdainful
Attitudes	Friendly	Aggressive
	Supportive	
		Dominant
	Attentive	Inattentive
		Gossip

Table 2: Recruiter affects

The computation of moods is based on emotions following ALMA [15]. In the context of our interview simulation, the period is determined by the number of cycle question/answer. Each answer leads to the computation of a new emotions set and these emotions influence the interviewer's mood. The basis for the calibration is as follows: after five cycles of a specific emotion (anger for example), the virtual

recruiter will be in the corresponding mood (hostile). More details about the mood computation can be found in [19].

The way we compute attitudes follow this principle: an agent can adopt an attitude according to its personality [33] or according to its actual mood [39]. For example, an agent with a non-aggressive personality may still show an aggressive attitude if its mood becomes very hostile. The mood compensates the personality and vice versa. For this reason, we use a logical-OR as condition on these two dimensions. As a consequence, in our model, the attitude can be triggered by one of these two dimensions. Then, the maximum value (mood or personality) is kept to compute the corresponding attitude, as is classically done in Fuzzy logics.

6.2 Decisional Core

Our main objective is to deduce real user beliefs from the *real user/virtual agent* interaction. In this interaction, inputs (affective states of the user) are given by non verbal signals deduced from social signal interpretation. It can be used on different simulations involving the interaction of a human with a virtual agent: teaching, formation, training, amongst other. The common feature of these simulations is the use of questions by the agent. Our model considers the use of questions in order to manage the context of the answers of the person interacting with the TARDIS.

To summarize, our theory of mind model has three main properties:

- The theory of mind is about a real person who interacts with the system
- It is centred on affective states interpretation of the person in front of the simulation
- It uses the context of questions to analyse user responses.

6.2.1 Context management

Labels are given to the questions/sentences of the virtual character in order to interpret the answer/reaction of the human in terms of beliefs with respect to some topics. A list of topics can be constructed for each specific application. The set of topics set_{topic} contains N topics: $\{topic_1, topic_2, \dots, topic_N\}$. Each subject is application dependent and based on the domain of the simulation. A question is defined by 0 to n topics.

6.2.2 Building a Beliefs Model

In order to build beliefs about the human who interacts with the system, we consider the questions/sentences that were just expressed by the virtual agent (identified by labels about topics) and the quality of the answer of the human from an affective point of view (which is obtained by the social cue recognition module) and the performance index P_d . Based on this information, the agent updates its beliefs about the human on a particular subject. We denote the beliefs of the agent about the human as $B_{Human}(topic_i)$ for i in $\{1, \dots, N\}$.

According to the topic(s) raised by the question/remark of the agent, beliefs will be updated. In pursuance of building

the beliefs of the human, we consider its answer (perceived via SSI) and based on P_d decide the degree to which the youngster's answer can be categorised as positive (+1), negative (-1) or neutral (0).

Based on the human's average answer (P_d) and the topic tags of the question/remarks just posed/made by the agent, the beliefs can be computed. Updates of each belief is done with the following formula for each topic :

Algorithm 1 Beliefs computation

```

for  $topic_i \in set_{topic}$  do
   $B_{Human}(topic_i) \leftarrow B_{Human}(topic_i) + \alpha \times P_d$ 

```

α is a weight between 0 and 1 that can be altered if we want the recruiter beliefs about the human to evolve quickly ($\alpha = 1$) or not (α near of 0). It can rely on the personality of the agent. An impulsive agent has an α near the value of 1 and a moderate one close to 0.

6.2.3 Desires and goals

The desires are used to define the strategic intentions of the agent. Desires are denoted by $D(B_{Human}(topic_i))$, because desires in an interaction are about beliefs of the human on a particular topic. It corresponds to the beliefs that every person process during an interaction according to the reaction of the interlocutor. For instance $D(B_{John}(football))$ is the desire of the agent about John's knowledge in the *football* topic.

6.2.4 Dynamics of goals

Social attitudes used can be defined on Leary circumplex [23]. As shown by Leary, attitudes can be separated into two categories: the positive ones (friendly, supportive, etc.) and the negative ones (aggressive, dominant, etc.).

Based on these two types of attitudes, we define the Algorithm 2 in order to enable for the desires of the agent to be updated.

Algorithm 2 Desires computation

```

if ( $Attitude \in set(attitude_-)$ ) then
  for  $B_{Human}(topic) \in set_{topic}$  do
    if ( $P_d < 0$ ) then
       $D(topic) \leftarrow D(topic) + \alpha \times |P_d|$ 
    else
       $D(topic) \leftarrow D(topic) - \alpha \times |P_d|$ 
if ( $Attitude \in set(attitude_+)$ ) then
  for  $B_{Human}(topic) \in set_{topic}$  do
    if ( $P_d < 0$ ) then
       $D(topic) \leftarrow D(topic) - \alpha \times |P_d|$ 
    else
       $D(topic) \leftarrow D(topic) + \alpha \times |P_d|$ 

```

This algorithm works as follows: if the agent has a negative attitude, its will intend to select topics with a negative answer for the human. On the other hand, if the agent has a positive attitude, its desires are about topics with a positive answer from the human.

6.2.5 Goal selection

Several strategies can be defined to enable the selection of one of the desires in the list of possible desires. The simplest strategy is to select the desire with the maximum value in the available desires. At any given point in the interview dialogue, every possibilities (topics) will not be suitable for preserving a logical sequence of the conversation (a scenario for instance).

7. CONCLUSION

In this article, we propose a pipeline to compute the affective reactions in a virtual recruiter from user behaviour. To this end, we employ the use of the SSI framework[38] to recognize user social cue in real time. The social cues are then forwarded to a online user model which uses them to compute a performance index of the user's behaviour. Based on the performance index as well as on the scenario, affects of the virtual recruiter are computed and a ToM approach allows to reason about the user's preferences.

Within the TARDIS training system, which is meant to aid young job seekers in acquiring job interview pertinent social skills, the pipeline will allow the virtual recruiter to react and adapt to the user's behaviour in real time, thus generating credible interaction.

As part of our future work, we aim to conduct studies to evaluate and validate the functionality of the described methodologies. On the technical side, we aim to improve the robustness of the pipeline, for example, by implementing mechanisms to explicitly deal with recognition errors. Furthermore, we plan to extend the online user model to compute other behavioural characteristics as well as to handle more social cues, such as expressivity features, physiological cues or eye gaze.

Acknowledgment

The research leading to this paper has received funding from the European Union Information Society and Media Seventh Framework Programme FP7-ICT-2011-7 under grant agreement 288578.

8. REFERENCES

- [1] K Anderson, E André, T Baur, S Bernardini, M Chollet, E Chrysafidou, I Damian, C Ennis, A Egges, P Gebhard, Hazaël Jones, Magalie Ochs, Catherine Pelachaud, K Porayska-Pomsta, Rizzo Paola, and Nicolas Sabouret. The TARDIS framework: intelligent virtual agents for social coaching in job interviews. *Proceedings of the Tenth International Conference on Advances in Computer Entertainment Technology (ACE-13)*. Enschede, the Netherlands. LNCS 8253, 2013.
- [2] Richard D. Arvey and James E. Campion. The employment interview: A summary and review of recent research. *Personnel Psychology*, 35(2):281–322, 1982.
- [3] Ruth Aylett, Ana Paiva, Joao Dias, Lynne Hall, and Sarah Woods. Affective agents for education against bullying. In *Affective Information Processing*, pages 75–90. Springer, 2009.

- [4] Simon Baron-Cohen. *Mindblindness*. MIT Press, Cambridge, Massachusetts, 1995.
- [5] Ligia Batrinca, Giota Stratou, Ari Shapiro, Louis-Philippe Morency, and Stefan Scherer. Cicero - towards a multimodal virtual audience platform for public speaking training. In Ruth Aylett, Brigitte Krenn, Catherine Pelachaud, and Hiroshi Shimodaira, editors, *Intelligent Virtual Agents*, volume 8108 of *Lecture Notes in Computer Science*, pages 116–128. Springer Berlin Heidelberg, 2013.
- [6] Tobias Baur, Ionut Damian, Patrick Gebhard, Kaśka Porayska-Pomsta, and Elisabeth André. A job interview simulation: Social cue-based interaction with a virtual character. In *Proceedings of the IEEE/ASE International Conference on Social Computing (SocialCom 2013)*, pages 220–227. IEEE, Washington D.C., USA, 2013.
- [7] Tobias Baur, Ionut Damian, Florian Lingenfelser, Johannes Wagner, and Elisabeth André. Nova: Automated analysis of nonverbal signals in social interactions. In *Human Behavior Understanding*, pages 160–171. Springer, 2013.
- [8] Paul Boersma and David Weenink. Praat: doing phonetics by computer [computer program]. version 5.3.56, retrieved 15 september 2013, 2013.
- [9] Tibor Bosse, Zulfiqar A Memon, and Jan Treur. A two-level BDI-agent model for theory of mind and its use in social manipulation. In *In AISB 2007 Workshop on Mindful Environments*, pages 335–342, 2007.
- [10] John Bynner and Samantha Parsons. Social Exclusion and the Transition from School to Work: The Case of Young People Not in Education, Employment, or Training (NEET). *Journal of Vocational Behavior*, 60(2):289–309, April 2002.
- [11] Jared Curhan and Alex Pentland. Thin slices of negotiation: predicting outcomes from conversational dynamics within the first 5 minutes. *Journal of Applied Psychology*, 92(3):802–811, 2007.
- [12] Ionut Damian, Tobias Baur, and Elisabeth André. Investigating social cue-based interaction in digital learning games. In *Proc. of the 8th International Conference on the Foundations of Digital Games*. SASDG, 2013.
- [13] Mehdi Dastani and Emiliano Lorini. A logic of emotions: from appraisal to coping. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 1133–1140, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
- [14] Florian Eyben, Martin Wöllmer, and Björn Schuller. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the international conference on Multimedia*, MM '10, pages 1459–1462, New York, NY, USA, 2010. ACM.
- [15] Patrick Gebhard. ALMA - A Layered Model of Affect. *Artificial Intelligence*, pages 0–7, 2005.
- [16] Daniel Goleman. *Social intelligence: The new science of human relationships*. Bantam, 2006.
- [17] Shlomo Hareli and Ursula Hess. What emotional reactions can tell us about the nature of others: An appraisal perspective on person perception. *Cognition & Emotion*, 24(1):128–140, 2010.
- [18] Mohammed Ehsan Hoque, Matthieu Courgeon, J Martin, Bilge Mutlu, and Rosalind W Picard. Mach: My automated conversation coach. In *International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2013)*, 2013.
- [19] Hazaël Jones and Nicolas Sabouret. TARDIS - A simulation platform with an affective virtual recruiter for job interviews. In *IDGEI (Intelligent Digital Games for Empowerment and Inclusion)*, 2013.
- [20] Ashish Kapoor and Rosalind W. Picard. Multimodal affect recognition in learning environments. In *Proceedings of the 13th annual ACM international conference on Multimedia*, MULTIMEDIA '05, pages 677–682, New York, NY, USA, 2005. ACM.
- [21] Felix Kistler, Birgit Endrass, Ionut Damian, Chi T Dang, and Elisabeth André. Natural interaction with culturally adaptive virtual characters. *Journal on Multimodal User Interfaces*, 6:39–47, 2012.
- [22] Andrea Kleinsmith and Nadia Bianchi-Berthouze. Form as a cue in the automatic recognition of non-acted affective body expressions. In *Proceedings of the 4th international conference on Affective computing and intelligent interaction - Volume Part I*, ACII'11, pages 155–164, Berlin, Heidelberg, 2011. Springer-Verlag.
- [23] T Leary. Interpersonal circumplex. *Journal of Personality Assessment*, 66(2):301–307, 1996.
- [24] Alan M Leslie. ToMM, ToBy, and agency: Core architecture and domain specificity. In L A Hirschfeld and S A Gelman, editors, *Mapping the mind Domain specificity in cognition and culture*, chapter 5, pages 119–148. Cambridge University Press, 1994.
- [25] Stacy Marsella, Jonathan Gratch, and J Rickel. Expressive behaviors for virtual worlds. *Lifelike Characters Tools Affective Functions and Applications*, pages 317–360, 2003.
- [26] A N Meltzoff. Understanding the Intentions of Others: Re-Enactment of Intended Acts by 18-Month-Old Children. *Developmental Psychology*, 31(5):838–850, 1995.
- [27] Magalie Ochs, Nicolas Sabouret, and Vincent Corruble. Simulation of the Dynamics of Non-Player Characters' Emotions and Social Relations in Games. *IEEE Transactions on Computational Intelligence and AI in Games*, 1:4:281–297, 2010.
- [28] Andrew Ortony, Gerald L Clore, and Allan Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, July 1988.
- [29] Helmut Prendinger and Mitsuru Ishizuka. Social role awareness in animated agents. *Proceedings of the fifth international conference on Autonomous agents AGENTS 01*, pages 270–277, 2001.
- [30] Helmut Prendinger and Mitsuru Ishizuka. The empathic companion: A character-based interface that addresses users' affective states. *Applied Artificial Intelligence*, 19(3-4):267–285, 2005.
- [31] David V Pynadath and Stacy C Marsella. PsychSim : Modeling Theory of Mind with Decision-Theoretic Agents. *Information Sciences*, 19(1):1181–1186, 2005.
- [32] Monika Sieverding. 'Be Cool!': Emotional costs of hiding feelings in a job interview. *International*

Journal of Selection and Assessment, 17(4):391–401, 2009.

- [33] Mark Snyder. The influence of individuals on situations: Implications for understanding the links between personality and social behavior. *Journal of Personality*, 51(3):497–516, 1983.
- [34] Andrea Tartaro and Justine Cassell. Playing with virtual peers: bootstrapping contingent discourse in children with autism. In *Proceedings of the 8th international conference on International conference for the learning sciences-Volume 2*, pages 382–389. International Society of the Learning Sciences, 2008.
- [35] Thomas V. McGovern, Barbara W. Jones, and Susan E. Morris. Comparison of professional versus student ratings of job interviewee behavior. *Journal of Counseling Psychology*, 26(2), 1979.
- [36] A. Vinciarelli, M. Pantic, D. Heylen, C. Pelachaud, I. Poggi, F. D’Errico, and M. Schroeder. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *Affective Computing, IEEE Transactions on*, 3(1):69–87, jan.-march 2012.
- [37] T. Vogt and E. Andre. Comparing feature sets for acted and spontaneous speech in view of automatic emotion recognition. In *Multimedia and Expo, 2005. ICME 2005. IEEE International Conference on*, pages 474–477, july 2005.
- [38] Johannes Wagner, Florian Lingenfelser, Tobias Baur, Ionut Damian, Felix Kistler, and Elisabeth André. The social signal interpretation (SSI) framework - multimodal signal processing and recognition in real-time. In *Proceedings of the 21st ACM International Conference on Multimedia, 21-25 October 2013, Barcelona, Spain.*, 2013.
- [39] Duane T Wegener, Richard E Petty, and David J Klein. Effects of mood on high elaboration attitude change: The mediating role of likelihood judgments. *European Journal of Social Psychology*, 24(1):25–43, 1994.
- [40] Zhihong Zeng, M. Pantic, G.I. Roisman, and T.S. Huang. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(1):39–58, jan. 2009.