Towards a Computational Architecture of Dyadic Rapport Management for Virtual Agents

Alexandros Papangelis, Ran Zhao, and Justine Cassell

ArticuLab, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, USA

Abstract. Rapport has been identified as an important factor in human task performance. Motivated by the proliferation of virtual agents that assist humans on various tasks, we propose a computational architecture for virtual agents, building on our own work on a dyadic model of rapport between humans and virtual agents. We show how such a system can be trained in order to build, maintain and destroy rapport.

1 Introduction and Related Work

While rapport, or feeling "in sync" with a partner, has sound theoretical foundations and demonstrated benefits in a variety of contexts, little success has been achieved in long-term rapport between a human and virtual agent / embodied conversational agent (ECA). To fill this gap, we analyzed existing social science literature, as well as our own data, and proposed a theoretical framework and computational model of rapport for human to virtual agent interaction (published in this same volume) [11]. Building on that work, we here propose a computational architecture that treats rapport as a dyadic phenomenon, and allows the virtual agent to manage it in real-time with human users. We argue that the proposed architecture enables the system to build, maintain and even destroy rapport, over multiple interactions with the same user. Our contribution, therefore, in this work is a computational architecture for real-time rapport management in human agent interaction, built on the dyadic model proposed in [11].

Some relational agents are designed for building long term social companionship, for example [1, 5, 6]. [3, 9] inter alia, propose agents that interact through verbal and non-verbal signals. [4] propose an architecture for generating social behavior in human to robot interactions, [12] takes achievement of synchrony into account, but the only other demonstration of rapport management comes from the VH Toolkit [2], which focuses only on non-verbal behavior, and only for "instant rapport." Our work improves on prior approaches by relying on strong theoretical foundations in the social sciences, as well as an analysis of peer tutoring data, which together allow us to construct a dyadic framework based on actual conversational strategies which are carried by specific behaviors and which achieve specific rapport goals.

T. Bickmore et al. (Eds.): IVA 2014, LNAI 8637, pp. 320–324, 2014.

2 Towards a Computational Architecture of Rapport

The dyadic nature of our architecture means that updates and grounding are achieved by taking into account both user and system state. More specifically, we follow [11] who define rapport-management strategies whose effect cannot be grounded until we observe the user's reaction. To achieve this it is necessary to represent a dyadic state that models what has been grounded; a model of the user representing the system's beliefs about the user; and a putative ECA state inside that user model, representing the system's beliefs of how the user perceives it. The proposed architecture is presented in Figure 1, where we extend a generic ECA architecture. Clear components denote generic modules, while shaded components denote our contribution to the overall architecture, where Intention Understanding interprets and then maps the intentions behind the user's actions to our model of rapport, the Friendship Classifier implements [10], Rapport carries out updates to the rapport state, and the Conversation Manager plans verbal and non-verbal output.

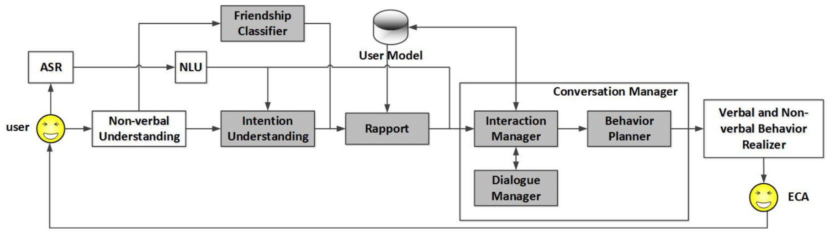

Fig. 1. The proposed VA architecture, incorporating our model of rapport

The most important data structures in our architecture, derived from our model, include the *dyadic state* representing the current state of rapport and a *user model* containing information acquired during the interaction. More specifically, the *dyadic state* comprises the following: 1) The System's goals, represented as a tree and split into task-oriented and social; 2) Rapport State, containing information about coordination, mutual attentiveness [8] and face [7]; 3) a behavioral model representing sociocultural and interpersonal norms (i.e. the ECA's behavioral expectations); 4) Friendship status, as a binary variable; and 5) History, containing information user act intentions, an estimate of the whether rapport is increasing or not, etc. The *user model*, contains information about the user's goals, shared knowledge, which may either be short-term (i.e. relevant to the current interaction only) or long-term, and a task model representing the user's progress regarding the task. In the *user model* we also represent a putative ECA state – an estimate of how the user perceives the system and comprising a rapport state, friendship status, shared knowledge and action intentions of the ECA.

In order to manipulate rapport a set of strategies have been defined, following [11], that allow the ECA to build, maintain and destroy rapport. For each strategy, a set of available system actions (or dialogue moves) $A_s \subset A$ can be defined, where A is the set of all (non-)verbal system actions; it is up to the Dialogue Manager (DM) to select the most appropriate ones. Rapport strategy selection is facilitated by taking into account the *dyadic state* and the *user model*. The selected strategy is then forwarded to the DM, responsible for selecting a set of appropriate actions, by taking into account the *dialogue state* which contains task- and interaction-related information. To assess rapport and update the *rapport state*, we estimate the user's intentions and, mapping these to the model, we update the *rapport state* accordingly[1]. Reinforcement Learning (RL) is a good candidate for learning which strategy and action to follow, paired with good feature selection methods. The behavioral model and dialogue policies, initialized to reflect general sociocultural norms for behavior in the particular context, are updated after each user action, according to how well the system's goals were met. As system and user interact, the strategy and dialogue policies gradually shift to reflect the increasingly interpersonal norms they follow. The policies, therefore, can be thought of as the facilitators of the rapport model, as they select strategies and actions based on the rapport state and current interpersonal norms. To have a way of measuring strategy success and update the behavioral model accordingly, we make a prediction of how the user should react to the chosen strategy, based on the rapport model, including the output from the friendship classifier and the putative ECA state (e.g. a FTA may have a different effect on strangers vs. friends). In order to make a prediction, we utilize the dyadic nature of our theoretical model (i.e. the putative ECA state and the *dyadic state* that applies to both ECA and human) and take advantage of the current (learnt) strategy selection policy, substituting the user model with the putative ECA state.

There are two phases where the Rapport module is used: rapport generation and rapport assessment. In the generation phase, we decide which rapport strategy to follow, based on the current *dyadic state* and the *user model*, while in the assessment phase, we assess the impact of the chosen strategy on the *dyadic state*, according to our model. During the rapport assessment phase, we observe the user's action and infer the intentions behind it, again according to our model, and use these inferred intentions to update the rapport state. It should be noted here that the dimensions of the rapport state pertaining to the Tickle-Degnen & Rosenthal [8] model, reflect an overall assessment of the system's and user's respective attentiveness and coordination. The system's face as well as the user's face are updated separately, immediately after performing a system or user action.

3 Concluding Remarks

Aiming to reduce complexity and improve tractability, it may be a reasonable first approximation to assume that task goals are independent of social goals and social features (*dyadic state* and *user model* separated from task model). Simplifying further, we assume that the overall strategy can be decomposed into

[1] `http://tinyurl.com/dyadic-rapport`

task-related and social-interaction-related moves. Thus, instead of learning an overall dialogue policy that achieves all the goals, we can now learn a task-related dialogue policy and a social-interaction-related policy. The selected strategy is forwarded to the DM, where action selection occurs by taking both dialogue policies into account. This, however, raises many interesting challenges such as dealing with competing strategies or incompatible actions (w.r.t task and social interaction), or including a module to aggregate the two strategies into one single strategy where possible. A complete treatment of this issue is kept for future work. We plan to train the system using our peer tutoring data and data from a Wizard of Oz study we plan to conduct, to train a simulated user that will interact with the system. To achieve this, we will apply inverse RL to estimate a good reward function and then direct sparse RL methods to train the simulator, and allow it to replicate and generalize from the data. Initially, competing goals will be addressed by the DM or the behavior planner.

As we increasingly refine the computational model presented in [11], we will incrementally implement our computational architecture, starting with a virtual peer reciprocal tutoring application for algebra. Foreseeable challenges include recognition of the human users' rapport strategies (which may span several dialogue turns or interleave with other strategies) in order to correctly update our model as well as react in the appropriate fashion, and training our learning modules before deployment for evaluation to ensure that our rapport strategies achieve the rapport management they are designed to convey.

Acknowledgments. The authors would like to thank Zhou Yu for her contribution to this work, as well as the other students and staff of the ArticuLab. This work was partially supported by the R.K. Mellon Foundation, and NSF IIS.

References

1. Bickmore, T.W., Caruso, L., Clough-Gorr, K., Heeren, T.: "It's just like you talk to a friend" relational agents for older adults. Interacting with Computers 17(6), 711–735 (2005)
2. Hartholt, A., Traum, D., Marsella, S.C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.-P., Gratch, J.: All together now: Introducing the virtual human toolkit. In: Aylett, R., Krenn, B., Pelachaud, C., Shimodaira, H. (eds.) IVA 2013. LNCS, vol. 8108, pp. 368–381. Springer, Heidelberg (2013)
3. Hoque, M.E., Courgeon, M., Martin, J.-C., Mutlu, B., Picard, R.W.: Mach: My automated conversation coach (2013)
4. Huang, C.-M., Mutlu, B.: Robot behavior toolkit: generating effective social behaviors for robots. In: Proc. of the 7th ACM/IEEE Int. Conf. on Human-Robot Interaction, pp. 25–32. ACM (2012)
5. Khosla, R., Chu, M.-T., Kachouie, R., Yamada, K., Yamaguchi, T.: Embodying care in matilda: an affective communication robot for the elderly in australia (2012)
6. Sidner, C.: Engagement: Looking and not looking as evidence for disengagement. In: Workshop at HRI, vol. 12 (2012)

7. Spencer-Oatey, H.: (im) politeness, face and perceptions of rapport: unpackaging their bases and interrelationships (2005)
8. Tickle-Degnen, L., Rosenthal, R.: The nature of rapport and its nonverbal correlates. Psychological Inquiry 1(4), 285–293 (1990)
9. Wilks, Y., Jasiewicz, J.: Calonis: An artificial companion for the care of cognitively impaired patients. In: MEMCA-14. This Proc. (2014)
10. Yu, Z., Gerritsen, D., Ogan, A., Black, A., Cassell, J.: Automatic prediction of friendship via multi-model dyadic features (August 2013)
11. Zhao, R., Papangelis, A., Cassell, J.: Towards a dyadic computational model of rapport management for human-virtual agent interaction. In: Bickmore, T., Marsella, S., Sidner, C. (eds.) IVA 2014. LNCS (LNAI), vol. 8637, pp. 514–527. Springer, Heidelberg (2014)
12. Zwiers, J., van Welbergen, H., Reidsma, D.: Continuous interaction within the saiba framework. In: Vilhjálmsson, H.H., Kopp, S., Marsella, S., Thórisson, K.R. (eds.) IVA 2011. LNCS, vol. 6895, pp. 324–330. Springer, Heidelberg (2011)