

Automatic Assessment and Learning of Robot Social Abilities

Francesco Del Duetto
L-CAS, University of Lincoln
Lincoln, UK
fdelduchetto@lincoln.ac.uk

Paul Baxter
L-CAS, University of Lincoln
Lincoln, UK
pbaxter@lincoln.ac.uk

Marc Hanheide
L-CAS, University of Lincoln
Lincoln, UK
mhanheide@lincoln.ac.uk

CCS CONCEPTS

• **Human-centered computing** → *Interactive systems and tools*; • **Computing methodologies** → **Online learning settings**; *Reinforcement learning*; • **Computer systems organization** → *Robotic autonomy*.

KEYWORDS

User Engagement; Social Human-Robot Interactions; Reinforcement Learning; Long-Term Autonomy

ACM Reference Format:

Francesco Del Duetto, Paul Baxter, and Marc Hanheide. 2020. Automatic Assessment and Learning of Robot Social Abilities. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20 Companion)*, March 23–26, 2020, Cambridge, United Kingdom. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3371382.3377430>

1 INTRODUCTION & BACKGROUND

One of the key challenges of current state-of-the-art robotic deployments in public spaces, where the robot is supposed to interact with humans, is the generation of behaviors that are engaging for the users. Eliciting engagement during an interaction, and maintaining it after the initial phase of the interaction, is still an issue to be overcome. There is evidence that engagement in learning activities is higher in the presence of a robot, particularly if novel [1], but after the initial engagement state, long and non-interactive behaviors are detrimental to the continued engagement of the users [5, 16]. Overcoming this limitation requires to design robots with enhanced social abilities that go past monolithic behaviours and introduces *in-situ* learning and adaptation to the specific users and situations. To do so, the robot must have the ability to perceive the state of the humans participating in the interaction and use this feedback for the selection of its own actions over time [27].

The research project that informs the work presented here is a collaboration between the University of Lincoln and The Collection Museum¹ with the objective of deploying an autonomous robot to engage with the museum's visitors and inform them about the local archaeology. The robot started operating on-site in October 2018, in an ongoing deployment of more than 1 year to date.

Being the museum a public space openly accessible to anyone, the interactions between the robot and the visitors are unstructured, in

¹<https://www.thecollectionmuseum.com/>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HRI '20 Companion, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7057-8/20/03.

<https://doi.org/10.1145/3371382.3377430>

the sense that users are not instructed how to interact with the robot. It is the latter that has to obtain and maintain the human attention in order to complete its tasks. This particular setting requires for an automatic method of assessing the robot behavior during the interactions, without the need of resorting to user interviews or questionnaires, and provides us with the opportunity of exploiting the real-world interactions for the task of improving the robot's social abilities over the long-term deployment.

Our research aim is, therefore, to embed the robot planning into a *reinforcement learning* (RL) framework, where the robot tries to improve its behavior to maximize the engagement of the humans it is interacting with. To do so, we propose the use of a model of user engagement, which provides an indirect assessment of the robot's social capabilities, that provides the RL reward and drives the learning process.

1.1 Characterization of Engagement

The definition of engagement during human-robot interactions has not been clearly specified yet [9], although, it can be described as a process composed of four parts: point of engagement, sustained engagement, disengagement, and re-engagement [17]. For the task of engagement detection in social interactions, approaches in the literature can typically be clustered in two groups: works that rely on specific behavioral cues, like gaze [2, 13, 20], context [4, 12, 22, 23] and other human perceptual features (e.g. people pose and sound) [7, 10, 15, 24], to define a rule-based scheme or to learn data-driven models and the more recent works that, taking a more holistic approach, learn models directly from engagement estimation or through proxy metrics [21, 28].

Assuming that humans are naturally able to accurately assess engagement in interactions, we propose to leverage their intuitive evaluation of it from videos of interactions, rather than relying on one particular definition of engagement. With the human coded engagement values we will build our own dataset and use it to learn an engagement prediction model from raw video data, framing the problem as a more general recognition task. In doing so, we answer the research question: “*can we learn a model to accurately measure users engagement during real-world interactions leveraging the human coded estimations of engagement?*”

1.2 Generating Social Behaviours

Previous work has shown that it is possible to influence the human engagement level during a human-robot interaction by employing different robot behaviors. Sidner et al. [25] explored how the use of gazing and gestures affects positively the user perception of the robot, increasing their engagement. Similarly, Holroyd [10] defines policies with the goal of increasing user engagement and shows that the robot equipped with these policies is perceived to be more

human-like, to behave more fluently and that users reciprocate more robot cues.

Recent works aimed at learning these social behaviors typically use (Deep) RL techniques to exploit the real-world interaction experiences a robot can collect. Qureshi et al. [18, 19] proposed end-to-end models to teach a robot the most appropriate action for approaching humans and starting an interaction. The reward signal was triggered by successful/unsuccessful handshakes. Lathuilière et al. [11] uses Deep RL to learn a gaze policy from an intrinsic reward function based on the audiovisual position of people with respect to the robot camera field of view. Gao et al. [8] learns a robot policy for approaching groups of people by maximizing a group formation score and minimizing the displacement of other participants in the group when the robot approaches.

In this project, we use state-of-the-art Deep RL approaches in order to improve our robot’s social behavior. We employ the human engagement level as a robot internal reward to maximize. Here we attempt to address the research question: “*is the engagement value estimated during the interactions a sufficient feedback signal for driving the in-situ learning of the robot social behaviors?*”

1.3 Contributions to the Field of HRI

In terms of contributions, the main objectives of this research are:

- (a) the proposal of a robotic framework for long-term deployment of robots in public spaces to promote and facilitate studies in-the-wild;
- (b) an analysis of user engagement with our robot over long-term deployment data;
- (c) a ready-to-use regression model for real-time estimation of users engagement during in-the-wild deployments;
- (d) the validation of the use of the predicted scalar engagement values to automatically assess the robot behavior, allowing *in-situ* adaptation and learning.

In our past and current work we have contributed to the community with the points (a), (b) and (c) as outlined in Sec. 2 and 3. Contribution (d) is left for future works and analyzed in Section 4.

2 LONG-TERM DEPLOYMENT AND THE NEED FOR ADAPTATION

The initial work of this project consisted in deploying our autonomous robot in a public museum. At this stage, we developed the tools necessary for the correct autonomous operations, such as: navigation stack, task scheduling, behavior specification, management and users interface. In our previous work [5] we describe the tools implemented, in particular we want to highlight here the integration of the Petri Net Plans (PNP) formalism [29] into our framework for the goal of specifying the robot behaviors and allowing their adaptation (see Sec. 4).

The robot can initiate 3 different interactive tasks with the visitors: (1) give a short verbal description of an exhibit; (2) guide the visitors to an exhibit and then describe it; (3) perform a guided tour centred around a theme, initially describing the theme of the tour and then guiding the visitors to each stop sequentially providing a description. Each task consists in a scripted behaviour in which the robot would always execute the same actions in the same manner. Only in tasks (2) and (3) there is a branching point in the behavior

where, during the description action, the robot gives initially little information about the exhibit and successively asks the users if they want to know more. Furthermore, at any moment during the interaction the users can stop the tasks before its natural termination. The tasks have an average duration of 20 seconds, 2 minutes and 10 minutes respectively, if not stopped by the user before their end. In our study [5], we have analysed the duration of a total of 5232 tasks, started over a period of 103 days of operation. The data suggests that:

- there is high initial engagement with the robot;
- the engagement is more and more difficult to maintain as the interaction proceeds further in time.

3 CONTINUOUS ENGAGEMENT ASSESSMENT

In our work [6], we asked independent human coders to annotate a dataset with a continuous per-frame engagement value, using an approach similar to Tanaka et al. [26]. We collected a total of more than 9 hours of video annotations and report a moderate to strong average inter-rater agreement over different smoothing factors (0.56 to 0.72 Spearman correlation). The dataset, which we named TOur GUide ROBOT (TOGURO), is composed of videos collected from the robot’s own camera during the interactions with users in our long-term deployment. It features a diverse range of interactions both in terms of audience demographics (e.g. number of people, age and gender) and in terms of its dynamics.

We successively trained an end-to-end engagement regression model, on the TOGURO dataset, to predict a value of engagement $y \in [0, 1]$ for each second of the video feed from the robot camera. The model reports a Mean Squared Error (MSE) of 0.126 on test data, validating the usefulness of the method in providing a continuous assessment of the users’ engagement in our experimental scenario. Moreover, an assessment of the trained model in predicting the loss of engagement over the publicly available UE-HRI dataset [3] reports an area under the ROC of 0.88, evidencing its generalization capability over completely different HRI scenarios.

4 LEARNING FROM ENGAGEMENT ESTIMATIONS

In the future, we plan to close the loop between the user engagement detection and the learning of social robot behaviors.

Having established that our engagement model can accurately predict engagement values during real-world interactions, we plan to use this model to generate online assessments of the robot behavior. The output of the engagement regression model will form the robot reward function which guides the learning. Therefore, the robot will learn to select the actions that generate higher user engagement at each moment during the interaction. We expect that, as the learning continues, the robot will be more and more able to sustain longer interactions with the museum’s visitors.

The robot behaviors, specified as PNPs, can be translated into stochastic policies, similarly to [14]. Therefore, we will bootstrap the learning process from the behaviors already defined by us allowing to start doing exploration from our ongoing deployment. We plan to learn both task-specific behaviors, like planning the sequence of exhibits in the tour, and more contingent behaviors, like gazing policies.

REFERENCES

- [1] Paul Baxter, Francesco Del Duchetto, and Marc Hanheide. 2018. Engaging Learners in Dialogue Interactivity Development for Mobile Robots. In *International Conference on Educational Robotics (EDUROBOTICS)*.
- [2] Paul Baxter, James Kennedy, Anna-Lisa Vollmer, Joachim de Greeff, and Tony Belpaeme. 2014. Tracking gaze over time in HRI as a proxy for engagement and attribution of social agency. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14*. ACM Press, New York, New York, USA, 126–127. <https://doi.org/10.1145/2559636.2559829>
- [3] Atef Ben-Youssef, Chloé Clavel, Slim Essid, Miriam Bilac, Marine Chamoux, and Angelica Lim. 2017. UE-HRI: a new dataset for the study of user engagement in spontaneous human-robot interactions. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction - ICMI 2017*. ACM Press, New York, New York, USA, 464–472. <https://doi.org/10.1145/3136755.3136814>
- [4] G. Castellano, I. Leite, A. Pereira, C. Martinho, A. Paiva, and P. W. McOwan. 2012. Detecting Engagement in HRI: An Exploration of Social and Task-Based Context. In *2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing*. 421–428. <https://doi.org/10.1109/SocialCom-PASSAT.2012.51>
- [5] Francesco Del Duchetto, Paul Baxter, and Marc Hanheide. 2019. Lindsey the Tour Guide Robot - Usage Patterns in a Museum Long-Term Deployment. In *International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, New Delhi.
- [6] Francesco Del Duchetto, Paul Baxter, and Marc Hanheide. 2020. Are you still with me? Continuous Engagement Assessment from a Robot's Point of View. [arXiv:cs.RO/2001.03515](https://arxiv.org/abs/2001.03515)
- [7] Mary Ellen Foster, Andre Gaschler, and Manuel Giuliani. 2017. Automatically classifying user engagement for dynamic multi-party human-robot interaction. *International Journal of Social Robotics* 9, 5 (2017), 659–674.
- [8] Yuan Gao, Fangkai Yang, Martin Frisk, Daniel Hernandez, Christopher Peters, and Ginevra Castellano. 2019. Learning Socially Appropriate Robot Approaching Behavior Toward Groups using Deep Reinforcement Learning. In *International Conference on Robot & Human Interactive Communication (RO-MAN)*. IEEE, New Delhi.
- [9] Nadine Glas and Catherine Pelachaud. 2015. Definitions of Engagement in Human-Agent Interaction. In *International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE Press, Xi'an, China, 944–949.
- [10] Aaron Holroyd. 2011. Generating engagement behaviors in human-robot interaction. (2011).
- [11] Stéphane Lathuilière, Benoit Massé, Pablo Mesejo, and Radu Horaud. 2018. Deep Reinforcement Learning for Audio-Visual Gaze Control. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 1555–1562.
- [12] Iolanda Leite, Marissa McCoy, Daniel Ullman, Nicole Salomons, and Brian Scasellati. 2015. Comparing models of disengagement in individual and group interactions. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 99–105.
- [13] Séverin Lemaignan, Fernando Garcia, Alexis Jacq, and Pierre Dillenbourg. 2016. From real-time attention assessment to with-me-ness in human-robot interaction. In *The eleventh acm/ieee international conference on human robot interaction*. IEEE Press, 157–164.
- [14] Matteo Leonetti and Luca Iocchi. 2010. LearnPNP: A tool for learning agent behaviors. In *Robot Soccer World Cup*. Springer, 418–429.
- [15] Marek P Michalowski, Selma Sabanovic, and Reid Simmons. 2006. A spatial model of engagement for a social robot. In *9th IEEE International Workshop on Advanced Motion Control, 2006*. IEEE, 762–767.
- [16] Illah R Nourbakhsh, Clayton Kunz, and Thomas Willeke. 2003. The mobot museum robot installations: A five year experiment. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, Vol. 4. IEEE, 3636–3641.
- [17] Heather L O'Brien and Elaine G Toms. 2008. What is user engagement? A conceptual framework for defining user engagement with technology. *Journal of the American society for Information Science and Technology* 59, 6 (2008), 938–955.
- [18] Ahmed Hussain Qureshi, Yutaka Nakamura, Yuichiro Yoshikawa, and Hiroshi Ishiguro. 2016. Robot gains social intelligence through multimodal deep reinforcement learning. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 745–751.
- [19] Ahmed Hussain Qureshi, Yutaka Nakamura, Yuichiro Yoshikawa, and Hiroshi Ishiguro. 2017. Show, attend and interact: Perceivable human-robot social interaction through neural attention Q-network. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1639–1645.
- [20] Charles Rich, Brett Ponsler, Aaron Holroyd, and Candace L Sidner. 2010. Recognizing engagement in human-robot interaction. In *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 375–382.
- [21] Ognjen Rudovic, Yuria Utsumi, Jaeryoung Lee, Javier Hernandez, Eduardo Castelló Ferrer, Björn Schuller, and Rosalind W Picard. 2018. CultureNet: A deep learning approach for engagement intensity estimation from face images of children with autism. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 339–346.
- [22] Hanan Salam and Mohamed Chetouani. 2015. Engagement detection based on multi-party cues for human robot interaction. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 341–347.
- [23] Hanan Salam and Mohamed Chetouani. 2015. A multi-level context-based modeling of engagement in human-robot interaction. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, Vol. 3. IEEE, 1–6.
- [24] Candace L Sidner, Cory D Kidd, Christopher Lee, and Neal Lesh. 2004. Where to look: a study of human-robot engagement. In *Proceedings of the 9th international conference on Intelligent user interfaces*. ACM, 78–84.
- [25] Candace L Sidner, Christopher Lee, Cory D Kidd, Neal Lesh, and Charles Rich. 2005. Explorations in engagement for humans and robots. *Artificial Intelligence* 166, 1-2 (2005), 140–164.
- [26] Fumihide Tanaka, Aaron Cicourel, and Javier R. Movellan. 2007. Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences* 104, 46 (Nov. 2007), 17954–17958. <https://doi.org/10.1073/pnas.0707769104>
- [27] Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. 2016. Computational human-robot interaction. *Foundations and Trends® in Robotics* 4, 2-3 (2016), 105–223.
- [28] Hae Won Park, John Busche, Bjorn Schuller, Cynthia Breazeal, Rosalind W Picard, et al. 2019. Personalized Estimation of Engagement From Videos Using Active Learning With Deep Reinforcement Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.
- [29] Vittorio A Ziparo, Luca Iocchi, Daniele Nardi, Pier Francesco Palamara, and Hugo Costelha. 2008. Petri net plans: a formal model for representation and execution of multi-robot plans. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 79–86.