ICDL-EPIROB
2019 Oslo, Norway

APRIL

SECURE
Safety Enables Cooperation in Uncertain Robotic Environments

Workshop on Personal Robotics and Secure Human-Robot Collaboration

(Joint APRIL and SECURE ITN Workshop)

Program

Monday, 19th August

Ingeniørenes Hus Møtesenter, Kronprinsens gate 17, 0251 Oslo

09:00 Keynote Speaker: Emre Ugur, Bogazici University, Istanbul

Learning to IMAGINE and condition the Action Consequences in Robotic Manipulation

10:00 SECURE students' presentations:

- **Chandrakant Bothe:**
  Language Learning for Safety during Human-Robot Interaction
- **Alessandra Rossi:**
  Investigating human perceptions of trust in robots for safe HRI in home environments

10:30 Coffee break

11:00 SECURE students' presentations (continuation):

- **Phuong D.H. Nguyen:**
  Toward Robots with Peripersonal Space Representation for Adaptive Behaviors

- **Marie Charbonneau:**
  On improving the coping capacities of whole-body controllers for humanoid robots
- **Francois Foerster:**
  Investigation on the Neuronal Mechanisms Involved in Accessing Learned Information in the Human Brain
- **Mohammad Thabet:**
  Sample-efficient Deep Reinforcement Learning with Imaginary Rollouts for Human-Robot Interaction

12:00 APRIL students' presentations:

- **Pontus Loviken:**
  Fast Online Model Learning for Controlling Complex Real-World Robots
- **Mina Marmpena:**
  Robotic emotional body language generation with variational autoencoders

12:30 Lunch

13:30 Keynote Speaker: Michael Spranger, SONY CSL, Tokyo

Autonomous Meaning Creation

14:30 APRIL students' presentations (continuation):

- **Bahar Irfan:**
  Multi-modal Personalisation in Long-Term Human-Robot Interaction
- **Alexandre Antunes:**
  Hierarchical recurrent neural networks for action and language grounding

15:00 Coffee break

15:30 Poster Spotlight

16:00 Poster Session

- **Alexis Billier:**
  Design of a robotic hand using the push-pull cable technology
- **Oksana Hagen:**
  Automatic discovery of priors for robotic learning

- **Egor Lakomkin:**
  Spoken Emotion Recognition with Deep Neural Networks for Human-Robot Interaction
- **Gregorios Skaltsas:**
  Can Human Physiological Data be used to Safely Adapt Robot Behaviour
- **Mohammad Ali Zamani:**
  Safer Human-Robot Interaction: Spoken Language-modulated Actions using Deep Reinforcement Learning
- **Asta Danauskiene, Mauricio Machado, Pinar Boyraz:**
  Multimodal Human Intent Recognition System for Collaborative Human-Machine Interaction
- **Lukas Hindemith, Anna-Lisa Vollmer, Britta Wrede, Frank Joublin:**
  People use interaction patterns to teach robots

17:00 End of the Event

Organisers:

- Angelo Cangelosi, University of Manchester, UK (APRIL ITN)
- Stefan Wermter, University of Hamburg, Germany (SECURE ITN)
- Alban Lafraquiere, SoftBank Robotics Europe (APRIL ITN)
- Sven Magg, University of Hamburg, Germany (SECURE ITN)

Abstracts

Keynote Speakers:

Emre Ugur, Bogazici University, Istanbul (9:00)

Learning to IMAGINE and condition the Action Consequences in Robotic Manipulation

Predicting the consequences of one's own actions is an important requirement for safe human-robot collaboration and their application to personal robotics. Neurophysiological and behavioral data suggest that human brain benefits from internal forward models that continuously predict the outcomes of the generated motor commands for trajectory planning, movement control, and multi-step planning. First, I will present our recent extension of propagation networks that enable the robot to predict the effects of its actions in scenes containing articulated multi-part multi-objects. Belief Regulated Dual Propagation Networks (BRDPN) consists of two complementary components, a physics predictor and a belief regulator. While the former predicts the future states of the object(s) manipulated by the robot, the latter constantly corrects the robot's knowledge regarding the objects and their relations. Next, I will talk on our recent learning from demonstration framework that is based on Conditioned Neural Processes. CNMPs extract the prior knowledge directly from the training data by sampling observations from it, and uses it to predict a conditional distribution over any other target points. CNMPs specifically learns complex temporal multi-modal sensorimotor relations in connection with external parameters and goals; produces movement trajectories in joint or task space; and executes these trajectories through a high-level feedback control loop. Conditioned with an external goal that is encoded in the sensorimotor space of the robot, predicted sensorimotor trajectory that is expected to be observed during the successful execution of the task is generated by the CNMP, and the corresponding motor commands are executed.

Michael Spranger, SONY CSL, Tokyo (13:30)

Autonomous Meaning Creation

In order to understand language, we need to understand how speakers of languages interact using language, how they acquire language and how they change and evolve the language they are using. In this talk I will discuss prior and recent experiments on the emergence of language within populations of robots. I will particularly talk about two types of experiments 1) experiments on language acquisition and 2) experiments on evolution of language. In both cases, we attempt to build social-interactionist systems with robots that mimic aspects of human communication. We are interested in algorithms that allow a learner to acquire a particular part of language from a tutor robot, e.g. spatial language, and/or develop a communication system completely from scratch. Importantly, in my work, language is always task specific and typically related to achieving particular goals in the real world (or in simulation). Here, compositional meaning is a crucial component as it allows speakers to convey new intentions easily. The talk will discuss how to model meaning, how meaning can be acquired and how it might evolve.

Students' presentations:

Chandrakant Bothe (10:00)

Language Learning for Safety during Human-Robot Interaction

Service robots need to show appropriate social behaviour in order to be deployed in social environments such as healthcare, education, retail, etc. Some of the main capabilities that robots should have are navigation and conversational skills. With the concern of safety during the human-robot verbal interaction, the aim is to study and research different linguistic aspects. In this work, we define the safety with the language such that the learned linguistic components could be used for inference for safer actions. For example, when a person gives sentiment cues from the utterance those could be used to understand the emotional state if the human user and also to learn safety cues [Bothe et al., 2017], [Lakomkin et al., 2017], [Bothe and Wermter, 2019]. Another example would be the politeness as a linguistic feature when a user is impatient, he/she uses smaller utterances like *go away* or *take me to some place* which is linguistically impolite, this gives a clue that the user wants a robot to navigate faster and vice versa [Bothe et al., 2018a]. Linguistic features that indicate politeness can provide social cues about a person's patient and impatient behaviour.

At the end of the work, we develop a dialogue system that incorporates linguistic features to provide safer robot actions. As language is very complex, different linguistic features could be used to assess the behaviour. In natural language understanding, dialogue act, which represents a functional type of utterance, plays a very important role in a dialogue system. We developed context-based neural models to recognize the dialogue acts [Bothe et al., 2018d], [Bothe et al., 2018b] and use them perform discourse analysis [Bothe et al., 2018c]. Results of research in this direction allow us to revisit the dialogue systems, develop and deploy on a robot to demonstrate a proof of concept where the linguistic features could be potentially used for safer human-robot interaction.

ACKNOWLEDGMENT

This project has received funding from the European Union's Horizon 2020 framework programme for research and innovation under the Marie Sklodowska-Curie Grant Agreement No. 642667 (SECURE).

REFERENCES

[Bothe et al., 2018a] Bothe, C., Garcia, F., Cruz-Maya, A., Pandey, A. K., and Wermter, S. (2018a). Towards dialogue-based navigation with multivariate adaptation driven by intention and politeness for social robots. In *International Conference on Social Robotics*.

[Bothe et al., 2017] Bothe, C., Magg, S., Weber, C., and Wermter, S. (2017). Dialogue-based neural learning to estimate the sentiment of a next upcoming utterance. In Lintas, A., Rovetta, S., Verschure, P. F., and Villa, A. E., editors, *Artificial Neural Networks and Machine Learning – ICANN 2017*, pages 477–485, Cham. Springer International Publishing.

[Bothe et al., 2018b] Bothe, C., Magg, S., Weber, C., and Wermter, S. (2018b). Conversational Analysis using Utterance-level Attention-based Bidirectional Recurrent Neural Networks. In *Proc. of the International Conference INTERSPEECH 2018*.

[Bothe et al., 2018c] Bothe, C., Magg, S., Weber, C., and Wermter, S. (2018c). Discourse-Wizard: Discovering Deep Discourse Structure in your Conversation with RNNs. *preprint arXiv:1806.11420*.

[Bothe et al., 2018d] Bothe, C., Weber, C., Magg, S., and Wermter, S. (2018d). A Context-based Approach for Dialogue Act Recognition using Simple Recurrent Neural Networks. In *Proc. of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, pages 1952–1957. European Language Resources Association (ERLA).

[Bothe and Wermter, 2019] Bothe, C. and Wermter, S. (2019). MoonGrad at SemEval-2019 Task 3: Ensemble BiRNNs for Contextual Emotion Detection in Dialogues. In *Proceedings of the International Workshop on Semantic Evaluation (SemEval-2019) at the Conference NAACL-HLT 2019*. ACL.

[Lakomkin et al., 2017] Lakomkin, E., Bothe, C., and Wermter, S. (2017). GradAscent at EmoInt-2017: Character and Word Level Recurrent Neural Network Models for Tweet Emotion Intensity Detection. In *Proc. of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis at the Conference EMNLP*, pages 169–174. ACL.

Alessandra Rossi, Kerstin Dautenhahn, Khang Lee Koay, Michael L. Walters **(10:15)**

Investigating human perceptions of trust in robots for safe HRI in home environments

Abstract—**Human-Robot Interaction (HRI) research has established that trust has a key role in human acceptance of robots in human-oriented environments. However, we can expect that robots will exhibit occasional mechanical, programming or functional errors, as with other electrical consumer devices. This research aims to identify the factors that undermine people's trust in robots, and to provide robots with coping mechanisms that will allow them to gain people's trust and to enhance a successful interaction. In this paper, we give a short overview of the state of our work.**

I. Introduction

The autonomous robots will enable to collaborate with people in their daily living activities in the future. This means that robots will need to be able to share the same physical space and engage human users in social interactions. Therefore, it is important that people accept and trust robots to be able to look after their well-being and assist them in a safe way. Trust can be affected by several factors [1] in HRI, which can be: Human-related, such as self-confidence and prior experience with robots; Robot-related, such as robot's embodiment and failure rates; and Environmental, such as communication and team collaboration. In this study, we explore some of these aspects which can affect and undermine human trust in robots. We provide strategies that can help the recovery of trust after a trust breach. In particular, we investigate the effects of a robot that exhibits social cues and that affects robot failures on people's trust in them. This research [2] investigate the following research questions:

RQ1 How do robot errors with varying magnitudes of consequences affect human trust in a robot?

RQ2 Does the impact on trust change if the error happens at the beginning or end of an interaction?

RQ3 Can the trust of humans in a robot be regained more easily if it is a big error happening at the beginning or end of an interaction? Or is it easier to recover from a loss of trust caused by a small error happening at the beginning or the end of the interaction?

RQ4 Does awareness of the robots' real potential and limitations affect human perceptions of trust in the robot?

RQ5 Are the use of human social behaviours by robots sufficient for humans to trust a robot to look after their well-being?

RQ6 Can a human's trust in their robot change over time (if the robot starts to show erratic behaviours)?

II. State of the Research

In this Section, we provide an overview of our current research findings:

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 642667 (Safety Enables Cooperation in Uncertain Robotic Environments - SECURE).

FQ1 People trust is affected more severely when the robot made errors having severe consequences ([3]–[5]).

FQ2 Participants tend to form their judgements at the beginning of an interaction and adjusted it later on, depending on the robot's performance ([3]–[5]).

FQ3 There is a greater tendency to not trust the robot when severe errors happen at the beginning of an interaction ([3]–[5]).

FQ4 For primary and secondary school students, we found that their awareness of being able to program different robots' behaviours affected their perception of the robots ([6], [7]).

FQ5 We conducted two different studies. One focused on a navigation task, and the other focused on tasks with different criticality levels. In the first study [8], we found that people preferred to interact with a social robot. They also tend to trust more a social robot to guide them than their initial expectations. In the second study, we did not observe differences between individual trust in a social or non-social robot.

FQ6 We are currently conducting a repeated study to assess how people's trust changes over several interactions. In particular, we aim to clarify research question RQ2.

References

[1] P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman, "A meta-analysis of factors affecting trust in human-robot interaction," *Human Factors: The Journal of Human Factors and Ergonomics Society*, vol. 53, no. 5, pp. 517–527, 2011.

[2] A. Rossi, K. Dautenhahn, K. L. Koay, and J. Saunders, "Investigating human perceptions of trust in robots for safe hri in home environments," in *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '17. New York, NY, USA: ACM, 2017, pp. 375–376.

[3] A. Rossi, K. Dautenhahn, K. L. Koay, and M. L. Walters, "Human perceptions of the severity of domestic robot errors," in *International Conference on Social Robotics (ICSR)*, K. Abderrahmane, Y. Eiichi, S. G. Shuzhi, S. Kenji, C. John-John, E. Friederike, and H. He, Eds. Tsukuba, Japan: Springer International Publishing, 2017, p. 647–656.

[4] ——, "A study on how the timing and magnitude of robot errors may influence people trust of robots in an emergency scenario," in *International Conference on Social Robotics (ICSR)*, K. Abderrahmane, Y. Eiichi, S. G. Shuzhi, S. Kenji, C. John-John, E. Friederike, and H. He, Eds. Tsukuba, Japan: Springer International Publishing, 2017, pp. 42–52.

[5] A. Rossi, K. Dautenhahn, K. Koay, and M. L. Walters, "The impact of peoples' personal dispositions and personalities on their trust of robots in an emergency scenario," vol. 9(1), pp. 137–154, 2018.

[6] A. Rossi, P. Holthaus, K. Dautenhahn, K. L. Koay, and M. L. Walters, "Getting to know pepper: Effects of people's awareness of a robot's capabilities on their trust in the robot," in *Proceedings of the 6th International Conference on Human-Agent Interaction*, ser. HAI '18. New York, NY, USA: ACM, 2018, pp. 246–252.

[7] A. Rossi, S. Moros, K. Dautenhahn, K. L. Koay, and M. L. Walters, "Getting to know kaspar: Effects of people's awareness of a robot's capabilities on their trust in the robot," in *submitted at IEEE International Conference on Robot and Human Interactive Communication 2019*, ser. RO-MAN 19, 2019.

[8] A. Rossi, F. Garcia, A. Cruz Maya, K. Dautenhahn, K. L. Koay, M. L. Walters, and A. K. Pandey, "Investigating the effects of social interactive behaviours of a robot on people's trust during a navigation task," in *to be published in Proceedings of the*, ser. TAROS 2019, 2019.

Phuong D.H. Nguyen, Matej Hoffmann, Ugo Pattacini, Giorgio Metta **(11:00)**

Toward Robots with Peripersonal Space Representation for Adaptive Behaviors

The abilities to adapt and act autonomously in an unstructured and human-oriented environment are necessarily vital for the next generation of robots, which aim to safely cooperate with humans. While this adaptability is natural and feasible for humans, it is still very complex and challenging for robots. Observations and findings from psychology and neuroscience in respect to the development of the human sensorimotor system can inform the development of novel approaches to adaptive robotics.

Among these is the formation of the representation of space closely surrounding the body, the Peripersonal Space (PPS), from multisensory sources like vision, hearing, touch and proprioception, which helps to facilitate human activities within their surroundings.

Taking inspiration from the virtual safety margin formed by the PPS representation in humans, this research first constructs an equivalent model of the safety zone for each body part of the iCub humanoid robot [1]. This PPS layer serves as a distributed collision predictor, which translates visually detected objects approaching a robot's body parts (e.g. arm, hand) into the probabilities of a collision between those objects and body parts. This leads to adaptive avoidance behaviors in the robot via an optimization-based reactive controller [2]. Notably, this visual reactive control pipeline can also seamlessly incorporate tactile input to guarantee safety in both *pre-* and *post*-collision phases in physical human-robot interaction (pHRI) [3]. Concurrently, the controller is also able to take into account multiple targets (of manipulation reaching tasks) generated by a multiple Cartesian point planner [4]. All components, namely the PPS, the multi-target motion planner (for manipulation reaching tasks), the reaching-with-avoidance controller and the human-centred visual perception, are combined harmoniously to form a hybrid control framework designed to provide safety for robots' interactions in a cluttered environment shared with human partners.

Later, motivated by the development of manipulation skills in infants, in which the multisensory integration is thought to play an important role, a learning framework is proposed to allow a robot to learn the processes of forming sensory representations, namely visuomotor [5] and visuotactile [6], from their own motor activities–motor babbling–in the environment. Both multisensory integration models are constructed with Deep neural networks (DNN) in such a way that their outputs are represented in motor space to facilitate the robot's subsequent actions.

ACKNOWLEDGMENT

Phuong D.H. Nguyen was supported by a Marie Curie Early Stage Researcher Fellowship (H2020-MSCA-ITA, SECURE 642667).

REFERENCES

[1] G. Metta, L. Natale, F. Nori, *et al.*, "The iCub humanoid robot: An open-systems platform for research in cognitive development," *Neural Networks*, vol. 23, no. 8-9, pp. 1125–1134, 2010.

[2] P. D. H. Nguyen, M. Hoffmann, A. Roncone, *et al.*, "Compact Real-time Avoidance on a Humanoid Robot for Human-robot Interaction," in *The 2018 ACM/IEEE Int. Conf. on Human-Robot Interaction*, ACM, 2018, pp. 416–424.

[3] P. D. H. Nguyen, F. Bottarel, U. Pattacini, *et al.*, "Merging physical and social interaction for effective human-robot collaboration," in *2018 IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, IEEE, 2018, pp. 710–717.

[4] P. D. H. Nguyen, M. Hoffmann, U. Pattacini, and G. Metta, "A fast heuristic Cartesian space motion planning algorithm for many-DoF robotic manipulators in dynamic environments," in *2016 IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, IEEE, 2016, pp. 884–891.

[5] P. D. H. Nguyen, T. Fischer, H. J. Chang, *et al.*, "Transferring Visuomotor Learning from Simulation to the Real World for Robotics Manipulation Tasks," in *2018 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, IEEE, 2018, pp. 6667–6674.

[6] P. D. H. Nguyen, M. Hoffmann, U. Pattacini, and G. Metta, "Reaching development through visuo-proprioceptive-tactile integration on a humanoid robot - a deep leanring approach," in *2019 Joint IEEE Int. Conf. on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*, IEEE, 2019.

Marie Charbonneau, Valerio Madogno, Luigi Penco, Francesco Nori, Serena Ivaldi, Daniele Pucci **(11:15):**

On improving the coping capacities of whole-body controllers for humanoid robots

I. Introduction

Within the next decades, humanoid robots are expected to autonomously and safely interact with complex dynamic environments and their inhabitants. Stable whole-body controllers are then a prerequisite to achieve safe robot motions. However, so far they can hardly cope with unexpected disturbances or with performing a variety of tasks. Developing controllers for *robustness* in these terms therefore appears to be fundamental, in order to increase autonomy.

This project explores methods to increase the robustness of humanoid whole-body controllers, along three major axes: joint limit avoidance, parameter tuning, and generalizing the whole-body motions that can be achieved by a controller.

II. Methods

An optimization-based whole-body torque-control framework based on the stack-of-tasks [1] is developed as the base of this work. It is formulated to be readily available for eventual dynamic locomotion tasks.

Controller parameters then need to be adjusted, in order to achieve an optimal robot behavior. We propose a framework to automatically learn task priorities, taking advantage of domain randomization [2] to encourage a higher robustness to external perturbations and different working conditions.

Nonetheless, optimal parameters may not always ensure joint limit avoidance, e.g. in case of external perturbations. As a solution, we propose novel feedback control laws based on parametrization of the joint space in terms of an activation function. It allows to cope with external perturbations, while convergence and stability can be theoretically proven.

Additionally, the effectiveness of a controller can be increased by achieving complex human-like whole-body motions. A whole-body teleoperation framework is then proposed, allowing the robot to keep balance while performing whole-body movements retargeted from a human operator.

III. Results

Results obtained with the whole-body torque-control framework show it to be suitable for balancing and walking of a humanoid robot. For instance, given reference trajectories, walking is achieved in [3]. Then, our method for learning task priorities has shown to be highly promising [4]. It automatically adjusts task priorities in simulation, while increasing the capability of the controller to cope with disturbances, changes in working conditions and platforms.

Joint limit avoiding control laws are proposed in [5]. They allow a torque-controlled robot to remain compliant, while resisting to external perturbations. Finally, a teleoperation control framework developed in [6] shows to allow real-time upper-body movements teleoperation, while the robot is simultaneously walking, taking footsteps adjusted according to those of the teleoperator.

IV. Conclusion

In summary, this project explores subjects related to whole-body control, and proposes methods that increase the capacity of humanoid robots to cope with disturbances and with performing a variety of motions. Exciting directions for future research may include (but are far from being limited to) testing joint limit avoidance in more challenging scenarios with the real robot, addressing the tuning of feedback gains and exploring real-time walking teleoperation. Eventually, exploring additional approaches that allow to increase even more the robustness achieved with whole-body controllers, promise to be highly beneficial. Not only will it save time and effort when deploying controllers, but it shall have a significant impact on the autonomy of robots within a human environment.

References

[1] N. Mansard, O. Stasse, P. Evrard, and A. Kheddar, "A versatile generalized inverted kinematics implementation for collaborative working humanoid robots: The stack of tasks," in *2009 International Conference on Advanced Robotics*, June 2009, pp. 1–6.

[2] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept 2017, pp. 23–30.

[3] S. Dafarra, G. Nava, M. Charbonneau, N. Guedelha, F. Andrade, S. Traversaro, L. Fiorio, F. Romano, F. Nori, G. Metta, and D. Pucci, "A control architecture with online predictive planning for position and torque controlled walking of humanoid robots," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2018, pp. 1–9.

[4] M. Charbonneau, V. Modugno, F. Nori, G. Oriolo, D. Pucci, and S. Ivaldi, "Learning robust task priorities of qp-based whole-body torque-controllers," in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, Nov 2018, pp. 1–9.

[5] M. Charbonneau, F. Nori, and D. Pucci, "On-line joint limit avoidance for torque controlled robots by joint space parametrization," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, Nov 2016, pp. 899–904.

[6] M. Charbonneau, L. Penco, F. Nori, D. Pucci, and S. Ivaldi, "A comprehensive framework for qp-based whole-body teleoperation," manuscript submitted for publication to *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

Francois Foerster (11:30):

Investigation on the Neuronal Mechanisms Involved in Accessing Learned Information in the Human Brain

New challenges in cognitive robotics relate to model cognitive abilities found in biological systems. A successful example of this inspiration is the powerful concept of affordances, that is the opportunities of action the environment offers to an animal [1], revolutionized cognitive robotics in by-passing the problem of storing knowledge of actions in artificial systems. The idea meant that perceptual inputs can be sufficient to estimate motor parameters [2], such as the extraction of the geometries of a visualized ground can inform whether it is safe to step on or how possibly maintaining balance on it.

However, studies on human cognition revealed that physical interactions with the environment cannot be based on sensorimotor loops solely [3]. For instance, manipulating an object for its function (e.g. tool use) necessitates the access to stored information abstracted over multiple sensorimotor experiences. Throughout experiences, cognitive agents learn how to interact with a given object in a safe manner: how to adjust a grip to not let it slip, what graspable part is likely harmless, and so on. Mimicking these abilities represent the most complex challenges that robotics in facing for decades.

Recent progress in cognitive neuroscience describes neural mechanisms involved in the processing of visual affordance [4] in humans and animals. However, what mechanisms the human brain uses to access to learned action and object representations remains unknown. Here, we suggest a theoretical framework allowing biological systems to access stored information represented in brain cortical areas during perception and interaction with objects in the environment. Because of their experience-dependent relationship, these object and action information are conceptualized as 'learned affordances' [5]–[7].

Using an original paradigm combining virtual reality coupled with electroencephalography (EEG), we report a series of five experiments recording the brain activities of ~160 human subjects. Results suggested that neuronal oscillations on a specific frequency band (at around 30 Hz) encodes the retrieval of object knowledge for action during perception. For instance, the function of an object, its name or how it is manipulated. More specifically, the automatic retrieval of these learned information has been mostly expressed as decreases of the signal power. This means that promising perspectives could be elaborated on these cortical signals to predict whether the brain accessed to stored information about the environment for adaptive behaviors in multiple situations, such as interacting with objects and possibly other humans.

We hope this investigation on the cognitive neuroscience of embodied memory and action systems in humans will lead to novel perspectives to build artificial systems able to interact safely with objects and humans.

REFERENCES

[1] J. J. Gibson, *The Ecological Approach to Visual Perception: Classic Edition*. 1979.

[2] G. Pezzulo and P. Cisek, "Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition," *Trends Cogn. Sci.*, vol. 20, no. 6, pp. 414–424, 2016.

[3] L. J. Buxbaum, M. F. Schwartz, and T. G. Carew, "The Role of Semantic Memory in Object Use," *Cogn. Neuropsychol.*, vol. 14, no. 2, pp. 219–254, 1997.

[4] P. Cisek and J. F. Kalaska, "Neural Mechanisms for Interacting with a World Full of Action Choices," *Annu. Rev. Neurosci.*, vol. 33, no. March, pp. 269–298, 2010.

[5] A. Antunes, G. Saponaro, A. Dehban, L. Jamone, R. Ventura, A. Bernardino, and J. Santos-victor, "Robotic tool use and problem solving based on probabilistic planning and learned affordances," pp. 11–13, 2015.

[6] M. A. Yasin, W. A. M. Al-Ashwal, A. M. Shire, S. A. Hamzah, and K. N. Ramli, "Denoising Auto-encoders for Learning of Objects and Tools Affordances in Continuous Space," *ARPN J. Eng. Appl. Sci.*, vol. 10, no. 19, pp. 8740–8744, 2015.

[7] L. Montesano, M. Lopes, and A. Bernardino, "Learning Object Affordances : From Sensory – Motor Coordination to Imitation," vol. 24, no. 1, pp. 15–26, 2008.

Mohammad Thabet (11:45):

Sample-efficient Deep Reinforcement Learning with Imaginary Rollouts for Human-Robot Interaction

Deep reinforcement learning (RL) has proven to be a great success in allowing agents to learn complex tasks from raw sensory information. Moreover, deep RL can be used to teach robots to interpret verbal and nonverbal cues from humans, which is essential if robots are to be used safely in social settings or in cooperative tasks with humans. However, the application of RL to actual robots can be prohibitively expensive, since it typically requires collecting huge amounts of interaction data with the environment. Furthermore, the unpredictability of human behavior in human-robot interaction tasks can hinder convergence to a good policy. Improving data efficiency for RL algorithms and the ability to learn in stochastic environments are therefore essential if RL is to be used efficiently in robots.

To this end, an architecture has been developed that allows agents to learn models of stochastic environments and use them to accelerate learning. Assuming that an environment model can be learned faster than an optimal policy, the agent can learn the environment model online simultaneously with the policy, and use the former to generate synthetic data to help learn the latter. The architecture consists of an encoder that compresses visual frames into low-dimensional abstract state representations, an environment model that is learned in this low-dimensional latent space, and a controller that maps these states into actions. Learning is performed from raw visual data, and the architecture assumes no prior knowledge of the environment, requiring only that the encoder be pretrained on task relevant images.

To validate the proposed architecture, an experiment has been designed in which a robotic arm has to solve a puzzle based on gestures made by a human. Results show that the proposed approach leads to significantly faster learning, and that the increase in performance due to the learned environment model is proportional to the difference in complexity between the environment dynamics and the task itself. The experiment also show that a learned environment model can be used to generate entire optimal plans, allowing the robot to solve tasks by observing just the initial state.

For future work, verbal cues will be included along with nonverbal cues. Moreover, a measure of model uncertainty will be used to limit the usage of synthetic data. Another important extension is to include recurrent models in the architecture to handle environments with non-Markovian states. Finally different ways of leveraging synthetic data will be investigated to improve data efficiency even further.

Pontus Loviken (12:00)

Fast Online Model Learning for Controlling Complex Real-World Robots

How can real robots with many degrees of freedom - without previous knowledge of themselves or their environment - act and use the resulting observations to efficiently develop the ability to generate a wide set of useful behaviours? This is an important question to the field of robotics, which in comparison to many other fields of computer science usually struggle with small datasets of samples, since acquiring the data is generally both expensive and time consuming [1-2]. So far however previous approaches are limited to cases with extensive pre-knowledge about the system, or where the robot has only a few degrees of freedom.

This work presents a novel framework that enables physical robots with many degrees of freedom to rapidly learn models for control from scratch. This can be done in previously inaccessible problem domains characterised by a lack of direct mappings from motor actions to outcomes, as well as state and action spaces too large for the full forward dynamics to be learned and used explicitly. The proposed framework is able to cope with these issues by the use of a set of local Goal Babbling models [3-5] that maps every outcome in a low dimensional task space to a specific action, together with a discrete higher level Reinforcement Learning model [6], that learns to navigate between the contexts from which each Goal Babbling model can be used. The two types of models can then be learned online an in parallel, using only the data a robot can collect by interacting with its environment.

To show the potential of the approach we present two possible implementations of the framework, over two separate robot platforms: a simulated planar arm with up to 1,000 degrees of freedom, and a real humanoid robot with 25 degrees of freedom. The results show that learning is rapid and essentially unaffected by the number of degrees of freedom of the robot, allowing for the generation of complex behaviours and skills after a relatively short training time. The planar arm is able to strategically plan series of motions in order to move its end-effector between any two parts of a crowded environment, within 10,000 iterations. The humanoid robot is able to freely transition between states such as lying on the back, belly, and sides, and occasionally also sitting up, within only 1,000 iterations. This corresponds to 30-60 minutes of real-world interactions.

The main contribution of this work is to provide a framework for solving a control learning problem, previously largely unexplored with no obvious solutions, but with strong analogies to, for example, early learning of body orientation control in infants. We examined two quite different implementations of the proposed framework, and showed success in both cases for two different control learning problem.

REFERENCES

[1] Mouret, J.B., 2016. Micro-data learning: The other end of the spectrum. *arXiv preprint arXiv:1610.00946*.
[2] Chatzilygeroudis, K., Vassiliades, V., Stulp, F., Calinon, S. and Mouret, J.B., 2018. A survey on policy search algorithms for learning robot controllers in a handful of trials. *arXiv preprint arXiv:1807.02303*.
[3] Rolf, M., Steil, J.J. and Gienger, M., 2011, August. Online goal babbling for rapid bootstrapping of inverse models in high dimensions. In *2011 IEEE International Conference on Development and Learning (ICDL)* (Vol. 2, pp. 1-8). IEEE.
[4] Benureau, F., 2015. *Self Exploration of Sensorimotor Spaces in Robots* (Doctoral dissertation).
[5] Baranes, A. and Oudeyer, P.Y., 2013. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, *61*(1), pp.49-73.
[6] Sutton, R.S. and Barto, A.G., *Reinforcement learning an introduction–second edition*, in progress, 2015.

Mina Marmpena, Angelica Lim, Torbjorn S. Dahl, Nikolas Hemion **(12:15)**

Robotic emotional body language generation with variational autoencoders

***Abstract*—Emotional body language is an important feature in human communication and expression. Humanoid robots in social environments can become more appealing and well accepted by using their embodiment to display cues of emotion expressions. We propose the application of deep learning methods, and more specifically the Variational Autoencoder model for the generation of emotional body language for the Pepper robot.**

I. Introduction

Emotional body language (EBL) is indispensable in human communication, hence its integration in robot behavior can enhance the user experience in human-robot interaction. So far it is a common practice to create robotic EBL animations with a hand-coded pose-to-pose process. This results in high quality animations, but limited in number and variability since it is a cumbersome process. We propose to use a generative model, the Variational Autoencoder that learns from a small set of high quality animations designed with the pose-to-pose method. In the sampling phase, the model can generate numerous new animations that appear smooth and realistic, while they exhibit increased expressive granularity.

II. Methods and Materials

A. Variational Autoencoders

The Variational Autoencoder (VAE) [1] is a probabilistic model that can be trained to encode the input into latent representations of lower dimensionality. These representations are stochastic and they are embedded into a continuous multivariate Gaussian manifold, the latent space, which can be sampled and decoded to generate new animations in the joint space of the robot.

B. Robot platform and training set

This study is applied on Pepper robot, a humanoid created by SoftBank Robotics. Pepper has 20 degrees of freedom. For this study we used 17 degrees of freedom, excluding the wheeled base. To train the VAE network, we used 36 animations designed by professional animators, with a pose-to-pose animation process. Essentially, each animation is a sequence of postures, each defined by the 17 angles. In total, the model was trained with 9392 postures. The animations have been previously validated in terms of readability and annotated with valence and arousal scores [2].

III. Results

A. Training

The model was trained for 200 epochs (reconstruction loss = 0.0028, variational loss = 4.2175). Fig.1 shows all the training set postures encoded in the 3D latent space and color-coded according to the animation they belong.

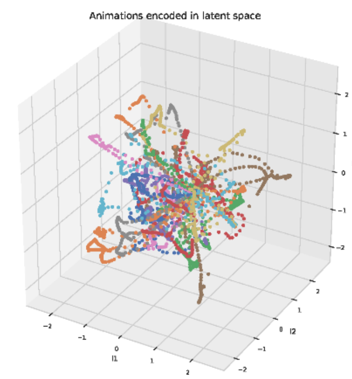

Fig. 1. The 3D latent space with the encoded animations.

B. Generating animations

For the generation of new animations we used different sampling and interpolation techniques. Besides, just randomly sampling the 3D latent space, interpolate the samples to get a latent interpolant, and then decoded it into the 17D joint space as an animation, we also applied spherical grid projections on the latent space to sample animations in a more systematic way, and explore for features that might be useful for more targeted emotion content [3].

References

[1] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv e-prints*, p. arXiv:1312.6114, Dec. 2013.

[2] M. Marmpena, A. Lim, and T. S. Dahl, "How does the robot feel? perception of valence and arousal in emotional body language," *Paladyn*, vol. 9, no. 1, pp. 168–182, 2018. [Online]. Available: https://doi.org/10.1515/pjbr-2018-0012

[3] M. Marmpena, A. Lim, T. S. Dahl, and N. Hemion, "Generating robotic emotional body language with variational autoencoders," in *Proceedings of the 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*. Cambridge, UK: IEEE Computer Society, 2019, accepted.

Bahar Irfan (14:30)

Multi-modal Personalisation in Long-Term Human-Robot Interaction

I. Introduction

Long-term human-robot interaction will be an integral part of domestic applications, rehabilitation and education in the future. However, long-term interactions based on a fixed set of behaviours, bring about certain challenges due to the repeated interactions with users, such as decreased user engagement. Personalisation can improve user engagement and facilitate building rapport and trust [1]. User recognition is an essential step towards achieving personalisation in the interaction for multiple users. Contrary to the common problems in the computer vision field, users might not be known in advance in a human-robot interaction scenario, which is known as an *open world recognition* problem. Moreover, users are encountered continuously and incrementally. Hence, deep learning approaches may fall short for meeting those requirements due to the *catastrophic forgetting* problem defined as the severe loss of performance on previously learned classes after the introduction of a new class [2]. Moreover, these approaches may take a vast amount of time for recognition and re-training in low computational power systems.

Incremental learning is not sufficient to adapt to the changes in the environment, such as those in the appearance of a user. *Online learning* allows for updating previous beliefs and user models to overcome this problem. Furthermore, inaccurate or noisy data, such as a blurry image or illumination changes, can result in low performance if only face recognition (FR) is used to identify users [3]. Ancillary physical or behavioural characteristics called *soft biometrics*, such as age and gender can be used to improve the recognition performance [4] and overcome issues related to similarities between users.

Consequently, we designed a multi-modal incremental Bayesian network (MMIBN) [5], which is the first in combining multi-modal biometric information for sequential and incremental learning for open world recognition.

II. Methodology

The proposed multi-modal incremental Bayesian network combines weighted soft biometrics, namely, gender, age, height and time of interaction, with a primary biometric, FR, through a naive Bayes classifier model.

We introduced a two-step ad hoc mechanism for MMIBN to identify if a user is known: (1) FR threshold is used to establish an "unknown" state, (2) the highest resulting posterior probability is compared to the second highest to evaluate the quality of estimation. When a new user is encountered, the Bayesian network is expanded to allow incremental learning by scaling the conditional probabilities of the modalities. Online learning is achieved through Expectation Maximization with an adaptive learning rate based on Maximum Likelihood.

III. Results

We initially evaluated our model with 14 participants over four weeks period [5]. We used the proprietary algorithms of the Pepper robot (SoftBank Robotics Europe) to obtain the multi-modal biometric information. Our results suggested that the proposed models increase the identification rate up to 1.4% in open-set and 4.4% in closed-set recognition compared to 90.3% of FR. Moreover, MMIBN performed worse with online learning. However, due to the low number of participants and the limited age range, we concluded that a larger dataset was necessary to evaluate the capabilities of the system.

Accordingly, we created a multi-modal long-term user recognition dataset with 200 users of varying characteristics based on the IMDB-Wiki dataset [6]. We compared our proposed approach to a state-of-the-art open world recognition approach, Extreme Value Machine [7] (EVM). The results showed that the proposed MMIBN models improved the identification rate significantly and substantially compared to soft biometrics, FR and EVM. Our results concluded that online learning either decreases recognition performance compared to fixed likelihoods or performs at the same significance level, which might result from accumulating noise of the identifiers.

IV. Ongoing Work

We applied the proposed user recognition model for a barista robot that learns the preferences of customers. We are currently analysing the results of a recent real-world study with 17 participants with repeated interactions over a week.

Moreover, we are using the proposed model for recognising users in cardiac rehabilitation therapy to personalise the interaction based on the information about the patients' previous sessions and their progress during the therapy [8]–[10].

References

[1] B. Irfan, A. Ramachandran, S. Spaulding, D. F. Glas, I. Leite, and K. L. Koay, "Personalization in long-term human-robot interaction," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2019.

[2] J. L. Mcclelland, B. L. Mcnaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory," *Psychological Review*, vol. 102, no. 3, 1995.

[3] W. Wójcik, K. Gromaszek, and M. Junisbekov, "Face recognition: Issues, methods and alternative applications," in *Face Recognition - Semisupervised Classification, Subspace Projection and Evaluation Methods*, S. Ramakrishnan, Ed. InTech, 2016, ch. 02.

[4] A. K. Jain, A. A. Ross, and K. Nandakumar, *Introduction to Biometrics*. Springer Publishing Company, Incorporated, 2011, ch. Multibiometrics.

[5] B. Irfan, N. Lyubova, M. Garcia Ortiz, and T. Belpaeme, "Multi-modal open-set person identification in HRI," in *2018 ACM/IEEE International Conference on Human-Robot Interaction Social Robots in the Wild workshop*, 2018.

[6] R. Rothe, R. Timofte, and L. Van Gool, "Deep expectation of real and apparent age from a single image without facial landmarks," *International Journal of Computer Vision (IJCV)*, 2016.

[7] E. M. Rudd, L. P. Jain, W. J. Scheirer, and T. E. Boult, "The extreme value machine," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 3, 2018.

[8] J. S. Lara, J. Casas, A. Aguirre, M. Munera, M. Rincon-Roncancio, B. Irfan, E. Senft, T. Belpaeme, and C. A. Cifuentes, "Human-robot sensor interface for cardiac rehabilitation," in *2017 International Conference on Rehabilitation Robotics (ICORR)*, 2017, pp. 1013–1018.

[9] J. Casas, B. Irfan, E. Senft, L. Gutiérrez, M. Rincon-Roncancio, M. Munera, T. Belpaeme, and C. A. Cifuentes, "Social assistive robot for cardiac rehabilitation: A pilot study with patients with angioplasty," in *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018.

[10] J. Casas, B. Irfan, E. Senft, L. Gutierrez, M. Rincon-Roncancio, M. Munera, T. Belpaeme, and C. A. Cifuentes, "Towards a SAR system for personalized cardiac rehabilitation: A patient with PCI," in *2018 ACM/IEEE International Conference on Human-Robot Interaction Personal Robots for Exercising and Coaching workshop*, 2018.

Alexandre Antunes (14:45)

Hierarchical recurrent neural networks for action and language grounding

Abstract—**A major aspect to make robots accessible to everyone in the future is the ability to interact using natural language. This is a major issue in the field of AI and robotics due to the complexity and sequential nature of grammar and language in general. In addition to this, the robot has to be able to link words and verbs to motor actions in the real world, thus bringing an additional level of complexity in action understanding and generation and the grounding of language with these actions. In this project we make the case for hierarchical recurrent neural network models for language generation and grounding, discussing the main advantages and shortcomings of such models when dealing with these tasks. We implement a Multiple Timescale Long Short-Term Memory (MT-LSTM), a hierarchical recurrent model that exhibits compositionality and generalisation properties, to connect language and motor actions, capable of generating a motor action from a sentence and vice-versa. We tested this model in both iCub and Pepper robots, analysing the model for these capabilities.**

I. INTRODUCTION

The future of robotics lies in close cooperation between humans and robots, which in turn requires robots to know and understand human language and to react correctly to it. While this is taken for granted with other humans who share a language and motor skills, it is still an open field of research in robotics. Robots need to learn language, connect it to their perception of the real world, and translate it into the appropriate motor actions that the human is expecting.

One way to address this issue is to look into the human brain, how it learns language and how it grounds this language in the world. Studies from neuroscience [1] highlighted a link between learning of motor actions and language, with a particular set of neurons activating for both motor action executions or sentences discussing this action. In studies regarding language, performed by Tomasello [2], there was strong evidence that children learn language by forming "verb-islands", generalising into a global grammar structure as more vocabulary was learned. These studies seem to suggest that language and actions are deeply connected, they are learned simultaneously, and both have sequential and compositional natures.

In this project we aim to recreate the human learning behaviour for actions and language, as a first step towards robots that can learn, and then act, as their human counterparts. One possible range of models that address the compositional and sequential natures of the task are hierarchical recurrent neural networks. A particular instance of these models, which has shown promising results on similar tasks, is the Multiple Timescales model, which consist of multiple layers of neurons, each with a different timescale factor. In previous works [3], [4] this model was used with a simple recurrent neural network, named Multiple Timescales Recurrent Neural Network (MTRNN), and was able to learn successfully learn both actions and language, albeit in separate models. This model serves as inspiration for the research presented in this project.

II. STATE OF THE RESEARCH

To implement a model that is able to learn language and actions simultaneously it needs to be able to generate both actions and sentences without any changes to the structure of the model, and should be trained in a bidirectional way.

Taking inspiration on the MTRNNs, we implemented a model that links an MTRNN used for action generation with one used for language generation. The two merged MTRNNs form a single model comprised of 7 layers with different timescales. The MTRNN model, however, proved too large to train, facing vanishing gradient problems. In order to deal with these problems, we adapted the model to use Long Short-Term Memory instead of simple Recurrent Neural Network cells. These MT-LSTM models have not been explored extensively so far, and the adaptation is not trivial or unique.

We implemented a 7 layer MT-LSTM model using the same premises as for the previous model. the MT-LSTM proved capable of learning to generate motor actions from a sentence and vice-versa, without changing the structure, being trained in a bidirectional way. Furthermore it still exhibits the compositional behaviour that is a requirement for this project. The model was tested in 3 different situations, with two different robots, iCub and Pepper.

One final question in the project is the following: How does this model compare to a child? Does it learn language similarly? Does it exhibit the "verb-island" behaviour discussed by Tomasello? Is it able to generalise a grammar structure like children do?

Some final experiments are being performed on the model that will answer these questions. We will investigate the training process, testing the network at different points and checking for evidence of verb-islands.

ACKNOWLEDGMENTS

This work is funded by the H2020 Marie Skłodowska-Curie Innovative Training Networks through the APRIL (H2020-MSCA-ITN-2015-674868) project.

REFERENCES

[1] Friedemann Pulvermüller and Luciano Fadiga. Active perception: sensorimotor circuits as a cortical basis for language. *Nature Reviews Neuroscience*, 11(5):351–360, May 2010.

[2] Michael Tomasello. The item-based nature of children's early syntactic development. *Trends in Cognitive Sciences*, 4(4):156–163, April 2000.

[3] Yuichi Yamashita and Jun Tani. Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS computational biology*, 4(11):e1000220, 2008.

[4] Stefan Heinrich, Cornelius Weber, and Stefan Wermter. Embodied language understanding with a multiple timescale recurrent neural network. In *International Conference on Artificial Neural Networks*, pages 216–223. Springer, 2013.

Posters:

Alexis Billier

Design of a robotic hand using the push-pull cable technology

***Abstract*— During this project, the main objectives were the creation of guidelines and the development of a system that will allow robots to utilize social cues that significantly improve safe human-robot interaction in home environments. This involves first investigating and understanding of 1) human safety risk during human-robot interaction in home environments, 2) the features that help humans to believe and trust that robots are able to look after their well-being, 3) social cues related to a human alert system that users are sensitive of and that can be implemented in robots to signal safety risks to the users. The project was divided in 2 different phase, the primary phase was dedicated to the development of a push-pull technology for the actuation of a robotic hand. The second phase was the implementation of this technology in the iCub robot during a secondment in IIT.**

I. Push Pull cable

The human hand is articulated by different tendons and muscles. For one finger, e.g. the index, there are several tendons which allow the movement of flexion and extension. Two for the flexion: one for the deep flexion and one for the superficial flexion. For the extension movement, more muscles and tendons are implicated in this motion. For the development of the humanoids robots, this architecture is often, especially in the cable-driven approach, where the motors and the cables replace respectively the muscles and the tendons. One of the main problems is the space required and the number of actuators used to move a finger. Most of the time at least two motors are used during the motion, meaning also two cables.

What this project investigated was to use the same cable for both flexion and extension. This way we divided by two the number of actuators and the cable. However, the main problem is the flexibility of the cable: it shall be flexible enough to go through the finger and the wrist and at the same time stiff enough to allow the extension.

Another problem that was investigated, is the lifting capacity of the robot. In fact, in the cable-driven approach, the lifting capacity is limited due to the architecture. The cable is directly connected to the motor, so to maintain the finger in opposition, the actuators have to keep the tension of the cable. This effort is supported by the motors, therefore the main limitation for the lifting is the stall torque of the motor. This paper proposes a new architecture of the mechanism which reduces the stall torque supported by the motor. The cables are not directly linked to the motor but to a trolley which can be locked in position, thus the tension of the cable is not supported by the motor.

II. Hand of iCub

The iCub hand has a lack of contact surface. It is due to the architecture of the finger, the actual architecture use cables to activate the fingers. As seen in the next figure the COR (Center of rotation) has great importance in this lack of contact surface.

To solve this problem a new architecture was investigated, combining both a push-pull cable technology and a bar linkage approach. It consists of a 3-link finger with two DoF. The finger consists of a crossbar system. The finger is only actuated in the Metacarpal-Proximal joint and the Proximal-Medial joint. The main advantage of this architecture is its simplicity. In fact, it uses only four parts to actuate the finger contrary to the previous architecture. In addition, the metacarpal-proximal joint and the proximal-medial joint are related to the displacement of the crank. Therefore, only one sensor is needed to know the exact position of the finger. This increases the repeatability of the finger.

As the hand size should be equivalent to the one of a three-year-old human hand, the space for the motors in the main hand body is limited. Therefore, the motors are located in the forearm of the robot. For the transmission of the movement between the motor and the finger, a cable used in push-pull configuration will be investigated.

In this way, only one cable per motor is needed for achieving both the flexion and the extension movements. Previous implementations required instead either two cables or one cable and passive return systems like spring for every actuated DOF.

The proposed system is the result of an optimization of the mechanical design to meet all specifications in terms of dimension, robustness, ability, and power. Five motors, one located in the palm, the other in the forearm, actuate the final hand. Fives cables transmit the movement between the motors and the fingers; these cables are used in both extension and flexion, with the use of the push-pull system.

This new architecture reduces the loss in the contact surface of 40%, with the previous architecture the loss of surface was $140mm^2$, the new loss is only $86mm^2$.

Oksana Hagen

Automatic discovery of priors for robotic learning

We strive for truly autonomous robots, that are able to adapt to new environments and learn new skills efficiently to be able to act in such scenarios as a care home or disaster rescue. Reinforcement learning techniques have demonstrated impressive performance, and have even exceeded human performance in such domains as ATARI games [1]. Good performance, however, comes at the cost of a very high sample inefficiency [4]. For example, RainbowDQN [6], the approach that features one of the best performances in the Atari games at the moment, passes the reference human performance after observing 18 million frames. This corresponds to about 83 hours of the play experience. If such training was performed with a real physical robot, it would inevitably lead to the tear of the robot or even breaking.

It is possible that his data-hungry behaviour of the state-of-art end-to-end reinforcement learning might be due to the fact that the systems always start de-nuovo and don't exploit any inherent regularities present in the environment. In our case, since we consider robotic learning, the tasks would always take place in the physical world, that has many unchangeable characteristics, such as continuity of time and space, laws of physics, etc. Our hypothesis is that such regularities could serve as powerful priors for increasing the data efficiency of task learning in robotics.

In this work we propose a closed-loop learning framework that allows extracting this prior knowledge on a pre-learning stage without any human input guided by the feedback from the performance of the agent on the task learning. The system consists of three optimisation loops, namely: task learning, state representation learning and learning of the loss function that is used for the state representation learning. This pre-learning procedure is performed on a set of randomised environments to ensure that the system can extract the common regularities of the world rather then particularities of a certain environment.

The experiments are performed on a Flatland simulated environment [5], developed by our lab. Flatland is a simple, lightweight environment for fast prototyping and testing of reinforcement learning agents. It is of lower complexity compared to similar 3D platforms (e.g. DeepMind Lab [2] or VizDoom [3]) but emulates physical properties of the real world, such as continuity, multi-modal partially-observable states with a first-person view and coherent physics.

References

[1] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning", *Nature*, vol. 518, no. 7540, p. 529, 2015.

[2] C. Beattie, J. Z. Leibo, D. Teplyashin, T. Ward, M. Wainwright, H. Küttler, A. Lefrancq, S. Green, V. Valdés, A. Sadik, J. Schrittwieser, K. Anderson, S. York, M. Cant, A. Cain, A. Bolton, S. Gaffney, H. King, D. Hassabis, S. Legg, and S. Petersen, "Deepmind lab", *CoRR*, vol. abs/1612.03801, 2016. arXiv: `1612.03801`. [Online]. Available: `http://arxiv.org/abs/1612.03801`.

[3] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, and W. Jaśkowski, "Vizdoom: A doom-based ai research platform for visual reinforcement learning", in *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, Sep. 2016, pp. 1–8. DOI: `10.1109/CIG.2016.7860433`.

[4] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people", *Behavioral and brain sciences*, vol. 40, 2017.

[5] H. Caselles-Dupré, L. Annabi, O. Hagen, M. Garcia-Ortiz, and D. Filliat, "Flatland: a Lightweight First-Person 2-D Environment for Reinforcement Learning", in *Workshop on Continual Unsupervised Sensorimotor Learning, ICDL-EpiRob 2018*, Tokyo, Japan, Sep. 2018. [Online]. Available: `https://hal.archives-ouvertes.fr/hal-01951945`.

[6] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M. Azar, and D. Silver, "Rainbow: Combining improvements in deep reinforcement learning", in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Egor Lakomkin

Spoken Emotion Recognition with Deep Neural Networks for Human-Robot Interaction

In the near future, robots will become common in home environments given recent advances in the fields of robotics, computer science, and machine learning. There are few challenges that are yet to be addressed to make robots our full-time companions: hardware limitations that currently restrict robots usage to laboratory conditions and software limitations that bias robots behavior to predefined scripts. Both of these affect the scalability of robots applications to real-life conditions. In this research, we focus on human-robot interaction, specifically on detecting potentially dangerous situations by analyzing acoustic information. For example, given an utterance identify if a person is excited or frustrated at the particular point of interaction. Speaker's affective state recognition is valuable information for developing human-robot interaction systems: for instance, robot's next action can be conditioned on the human's emotion state or detecting emotions like anger can be a clue for an unsafe situation.

I identify three main directions in my research: 1) features and signal representations learning for speech emotion recognition (SER) task. 2) investigation of neural architectures which allow robust to an internal robot's and an environmental noise emotion recognition 3) research on the methods and approaches to incorporate information contained in modalities other than auditory to improve speech emotion recognition. For example, linguistic analysis of a spoken text can help in difficult situations when analyzing only the acoustic signal is not enough to infer an affective state of the speaker.

I evaluate several dual architectures which integrate representations of the automatic speech recognition (ASR) neural network: a fine-tuning and a progressive network. The fine-tuning architecture reuses features learned by the recurrent layers of a speech recognition network and can use them directly for emotion classification. Our experiments on the IEMOCAP dataset show state-of-the-art performance in the accuracy and F1-score over the baseline recurrent neural network which is trained end-to-end for emotion recognition. Preliminary results were published and presented at the IJCNLP 2017 conference, followed by a journal submission with extended results and experiments. In addition, the importance of using speech recognition hypothesis for affective state recognition was presented at the ICRA 2019 conference.

Typically emotion recognition models start processing at the end of each utterance, which not only requires a mechanism to detect the end of an utterance but also makes it difficult to use them in a real-time communication scenario. We propose the EmoRL model that triggers an emotion classification as soon as it gains enough confidence while listening to a person speaking. As a result, we minimize the need for segmenting the audio signal for classification and achieve around 50% latency reduction as the audio signal is processed incrementally. The results will be presented at the ICRA 2018 conference (https://www.youtube.com/watch?v=js_TCxl_wF4).

We evaluate the robustness of state-of-the-art neural acoustic emotion recognition models in human-robot interaction scenarios. We observe a significant degradation in performance of neural models trained on datasets like IEMOCAP and afterwards evaluated on a real iCub Robot in the unobserved environment. We conduct several experiments and propose several ways to reduce the gap between the model's performance during training and testing in real-world conditions. The results were published at the IROS 2018 conference.

In this research, several methods and approaches were proposed and evaluated for robust emotion recognition for human-robot interaction. There are several challenges that are yet to be addressed: data efficiency of deep neural models, speech recognition in noisy far-field scenarios and learning emotion recognition on-line with partial feedback from human or sparse labels.

Gregorios Skaltsas

Can Human Physiological Data be used to Safely Adapt Robot Behaviour

a) Research and Professional Development

Attendance of SECURE project conferences and workshops except for one. University of Hertfordshire (UH) development seminars and tutorials were also attended. A two weeks summer school focused on advanced statistics and machine learning was attended during summer of 2016 in Madrid, Spain (http://www.dia.fi.upm.es/?q=es/node/275).

Two secondments took place. One where an industry-based human-robot interaction (HRI) experiment was setup and executed in the warehouse of OCADO (ocado.com) in Hatfield UK and the resulting recorded data was analyzed in Switzerland at École Polytechnique Fédérale de Lausanne (epfl.ch) as well as by remote collaboration. The second was the organization of an HRI experiment at Technical University of Vienna (tuwien.ac.at) which involved PEPER robot (softbankrobotics.com). Machine learning and statistics, networking and ethical hacking, sensor fusion (udemy.com) and free Matlab/ Python classes were also taken.

b) Research Topic Selection

The selection of the research topic combined data acquisition from multiple sources, data analysis, robot programming, human bio-signals. Methods that were recently and concurrently being investigated by researchers but without in-depth and conclusive findings were chosen and focus was given in the investigation of how the measurements may vary over time. Therefore, how habituation patterns are expressed in terms of stress response during HRI is one of the main research questions.

The title given so far is: "Can Human Physiological Data be used to Safely Adapt Robot Behaviour?". In my approach, one of the parameters the robot uses as feedback from the user is the user's stress levels. Stress detection is based on the user's physiological responses comprising mostly of galvanic skin response (GSR), heart-rate measurements and eye-tracking.

c) Experimental Setup and Results Analysis

There have been so far two experiments run. Both had as their main element the robot approaching the human participant in different speeds. This was implemented by following a trajectory that might be perceived as risky at some point for the task's successful execution or the robot's capability to localize or to stop at the right distance when reaching the human. The participants were split in groups that were experiencing four different scenarios in a different order in both experiments.

The analysis consists of analysis of the data obtained in relation to specific events. Standard signal processing rules and statistical processing are applied for the analysis of the results.

d) Current Findings and Scientific Publications

Current findings are under continuous revision and an upcoming publication is being prepared summarizing the findings of both experiments run. There seems to be a seasonal variation of the stress signals of the humans, as well as discreet variations in the perception of the robot's movement depending on the participants' acquaintance with the technology. Stress seems to be rising upon an event anticipation. However, it seems to be rising again if the intensity of the events is rising after a gradual drop.

Currently participated in the publication submitted to RoMAN 2019 titled "Evaluation of an Industrial Robotic Assistant in an Ecological Environment" which is related to the activities of my first secondment as described above.

e) Future Work

The future work planned is one or more small scale experiments to verify specific results, submission of the paper in preparation to a peer-reviewed journal or conference and the write-up of the thesis required for the completion of the PhD.

Mohammad Ali Zamani, Sven Magg, Cornelius Weber, Stefan Wermter

Safer Human-Robot Interaction: Spoken Language-modulated Actions using Deep Reinforcement Learning

In the future, robots are expected to work as companions with humans in various areas including domestic scenarios such as care-giving. However, even with well-engineered robots, it would be unrealistic to move robots directly from factories to home environments to perform complex tasks due to safety. Moreover, robots also have to continuously adapt to new environments to avoid hazardous actions since using experts to program a robot for every environment is impossible. Hence, we need adaptive learning algorithms.

Spoken language can be considered one of the most effective communication channels to warn robots about threats. A human can guide the robot by a verbal utterance toward a safer interaction. Our goal is to train a robot to safely perform complex tasks with the ability of processing environmental feedback, including guidance and warnings by a human, to shape a proper signal for updating its own policy. Therefore, this research is focused on three areas: a) given the verbal instructions, generating high-level actions, b) learning low-level actions to fulfill the high-level action, c) learning warnings using the prosodic/sentiment features of the human speaker.

A. From Spoken Instruction to a Sequence of Actions

We introduced a framework to obtain the intention of a given spoken instruction (e.g. "`boil water`") and generate the sequence of actions ("`moveto kettle`", "`grasp kettle`", ...) to fulfill the task [1], [2]. The intention detection was implemented with an MLP and trained on the TellMeDave corpus to predict one of 10 predefined classes.

We developed a symbolics environment from the "Tell Me Dave" Corpus to train the RL agent. The main contribution was to use a distributed compositional state representation (e.g. {`On Kettle Sink`}, ...) which reduced the learning time on given tasks. Our Reinforcement Learning (RL) successfully learns to perform the given task. Our RL is built based on Deep Q-Network with improvement to support multiple Q functions and different types of value estimation.

B. From high-level actions to low-level actions

Safety becomes more important when humans work with robots collaboratively especially in the fine-grained action level. For shaping such a safe collaborative scenario incrementally, as an initial step, we improved the learning of the Deep Deterministic Policy Gradient (DDPG) in a reach-for-grasp task by introducing an adaptive (larger-than-life) augmented target [3]. Later, we used it to train a 2-DOF arm in an interactive scenario to reach multiple target points which improved the learning time by solving the problem in simulation and deploying it on the robot when it gained enough confidence [4]. Then, we demonstrated how spoken instructions can be mapped to a spatial representation of the robot's workspace which can be used as constraints for the path planner [5].

C. Extracting Warning from the Human Speech

The robot needs to continuously process human speech to detect implicit interruptions or any change in the instruction. The robot is expected to be able to stop (both soft and emergency) with minimum latency in an unsafe situation.

We developed a model which was consist of Gated Recurrent Units (GRU) which learns a temporal representation from the extracted features of speech [6]. The Emotion classification module used the GRU's output to determine emotion as angry or neutral. The action selection module which is Monte Carlo Policy Gradient (or REINFORCE) decides to either wait for the next speech frame or terminate the processing and read the emotion classification module.

As a result, our model achieved about 50% latency reduction with the same level of accuracy evaluated on the iCub recorded data in our lab. We also improved the robustness of emotion recognition by proposing data augmentation techniques like overlaying background noise [7]. Moreover, by combining speech recognition with deep neural character-level language model, we achieved competitive performance for recognizing the speaker's sentiment [8].

D. Conclusions

In this PhD project, spoken instructions are used to instruct the robot. These instruction are mapped into a sequence of high-level actions depending on the state of the environment. Moreover, each high-level action is composed of fine-grained actions. The ability to modulate them depending on the situation leads to a safer human-robot collaboration which we demonstrated it for the reach-for-grasping task as an important primitive action. For detecting warnings, We also proposed a model to recognize emotions more robust and quicker, which can be used as an implicit interruption to planning to lead to a safer human-robot interaction.

References

[1] M. A. Zamani, S. Magg, C. Weber, and S. Wermter, "Deep reinforcement learning using symbolic representation for performing spoken language instructions." *In the BAILAR Workshop on Robot and Human Interactive Communication (RO-MAN)*, 2017.

[2] ——, "Deep reinforcement learning using compositional representations for performing instructions." *Paladyn Journal of Behavioral Robotics*, pp. 358–373, 2018.

[3] M. Kerzel, H. Beik Mohammadi, M. A. Zamani, and S. Wermter, "Accelerating deep continuous reinforcement learning through task simplification," *Int. Joint Conf. on Neural Networks (IJCNN)*, 2018.

[4] H. Beik Mohammadi, M. A. Zamani, M. Kerzel, and S. Wermter, "Online continuous deep reinforcement learning for a reach-to-grasp task in a mixed-reality environment," *accepted at International Conference on Artificial Neural Networks (ICANN)*, 2018.

[5] M. A. Zamani, H. Beik Mohammadi, M. Kerzel, S. Magg, and S. Wermter, "Learning Spatial Representation for Safe Human-Robot Collaboration in Joint Manual Tasks," *WORKMATE, ICRA*, 2018.

[6] E. Lakomkin, M. A. Zamani, C. Weber, S. Magg, and S. Wermter, "Emorl: Continuous acoustic emotion classification using deep reinforcement learning," *Int. Conf. on Robotics & Automation (ICRA)*, 2018.

[7] ——, "On the robustness of speech emotion recognition for human-robot interaction with deep neural networks," *International Conference on Intelligent Robots and Systems (IROS)*, pp. 854–860, 2018.

[8] ——, "Incorporating End-to-End Speech Recognition Models for Sentiment Analysis," *ICRA*, 2019.

Asta Danauskiene, Mauricio Machado, Pinar Boyraz

Multimodal Human Intent Recognition System for Collaborative Human-Machine Interaction

Chalmers University of Technology, Department of Mechanics and Maritime Sciences (M2),
412 96, Göteborg, Sweden

Abstract

Purpose:

This study focuses on developing an open source platform for multimodal human intent recognition investigating the electrical brain activity, eye gaze and pupil behavior in the scope of goal-directed movement intention recognition for human-machine interaction applications. Previous studies support that the electroencephalography (EEG) data is suitable for early motion recognition and prediction and the pupil size changes correlate with the difficulty of the task. However, few studies have investigated neural correlates of goal-directed and no-goal movements as well as the correlation between the pupil changes, EEG data and hand motion. We aim to uncover these correlations and obtain a multi-modal intent recognition system that can be utilized for collaborative tasks where human and robots interact for better performance in terms of speed and accuracy.

Methodology:

We designed a set of cue-based movement experiments that include (i) changing goal, (ii) repeating goal and (iii) no-goal scenarios which involve tracking tasks, performed in collaboration with a robot. The results were analyzed using multimodal data, regarding movement related cortical potentials (MRCP) and event related spectral perturbation (ERSP) of EEG data, evoked pupil response, eye-gaze movement patterns as well as binary goal\no-goal classification of the data and correlation between different bio-signals.

Results:

Our results indicate that changing goal-directed movements are distinguishable from no-goal movements in EEG data in both temporal and time-frequency domains, when performing the task with a passive robot. Collaborative robot experiments showed great inter-subject variability, therefore need to be further investigated. No correlation between evoked pupil response and MRCP was found in this study, however results suggest a correlation between MRCP and motion velocity profile.

Original Value:

This study proposes an open-source data collection method utilizing Lab Streaming Layer enabling multi-modal data collection that can be used to investigate human intent recognition. All the hardware utilized also feature open-source and modular structures such as open-source eye-tracking and equipment from Pupil-Labs and open-source BCI (brain-computer-interface). In addition to this contribution towards ***open-science***, connecting all state-of-the-art low-cost and open sources, the work involves preliminary investigation of automatic recognition of goal vs. non-goal actions of humans while collaborating with a 3DOF SCARA robot arm. The efforts spent in designing the intent recognition system also revealed weak or strong correlations between eye movement (i.e. gaze) patterns, EEG related data and action/motion metrics such as velocity. Although it cannot be firmly confirmed, we also observed and reported the effects of 'learning' in the action parameters in the repeated tasks which could be utilized for educational robotic tools to enhance learning in child development or rehabilitation of recovering adult patients after a neurological disruption to their motor-skills.

Keywords: Safe Human-Robot Interaction, Human Intent Recognition, Collaborative HMI, Goal-Directed Movement, Movement Prediction, Gaze Tracking, Pupillometry, BCI.

Lukas Hindemith[1], Anna-Lisa Vollmer[1], Britta Wrede[1], Frank Joublin[2]

People use interaction patterns to teach robots

[1]Applied Informatics - Bielefeld University, Bielefeld, Germany

[2]Honda Research Institute Europe, Offenbach, Germany

I. PURPOSE

Our work aims to make a first step towards the realization of *pragmatic frames* (PFs) for improved human-robot interaction (HRI). Pragmatic frames is a theoretical framework developed through observations in Parent-Child-Interaction. The observations showed that for learning situations parent and child co-construct a recurring flexible interaction protocol, which facilitates learning [1]. For example imagine a situation where a mother and her son are reading a picture book (cf. [5]). In this labeling frame, the mother points to an image and says "Look!" to direct the child attention. Then she asks her son "What's this?". He responds and she, independently of his performance, gives, mostly positive, feedback and provides the correct label. To transfer these observations to HRI, Vollmer et al. [2] analyzed the PFs used in current approaches to robot learning in interaction and found that current interaction protocols are artificial, mostly pre-programmed into the system and can't adapt to humans. To understand if and how PFs are used by non-experts when teaching a robot, we conducted a preliminary study. Based on the results, we specified requirements for the framework.

II. METHODOLOGY

In our study, eleven subjects were asked to teach the robot Pepper [3] how to cook, based on a given recipe as well as toy tools and ingredients with attached ArUco-markers [4]. Pepper was equipped with a set of reactive behaviors, but no learning abilities:

- Spotted keyword: repeat keyword
- Moving object detected: display object image on tablet
- Action keyword and moving tool: predefined arm-movements for action

The participants were instructed that in order for Pepper to be able to help somebody with cooking, it lacks some basic skills such as movements when stirring, cutting etc. and it neither knows the ingredients nor the utensils, which they all should teach the robot. The experiment took place in a kitchen, where the participant and Pepper stood opposite each other at a working surface. For the following analysis of the data we annotated the video-material regarding different dialog-acts, like speech or object manipulation.

This work was supported by Honda Research Institute Europe GmbH, Offenbach am Main, Germany.

III. RESULTS

Our analysis on labeling PFs showed that on the level of dialog-acts, one interaction pattern could be observed not only within each subject, but also between subjects. The structure of the pattern can be seen in Figure 1.

While this pattern – which we assume to be the deep syntax

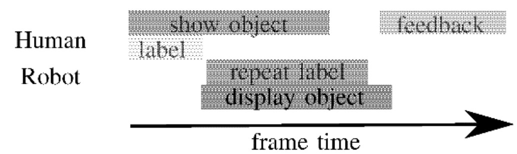

Fig. 1. Labeling PF

of the PF (cf. [1]) – can be found in most subjects, the concrete realization of each dialog-act differs. Analyzing speech and object manipulation qualitatively for two subjects revealed that on a micro-level, the behavior remains the same within a subject but differs greatly between subjects (holding the object and moving it towards the robot vs. placing the object on the table in front of the robot). Similar observations could be made with utterances introducing the object label. Additionally, PFs change on a micro-level within an interaction, which seems to come along with a changing mental model the subject has of the robot (e.g., at the beginning, one subject was happy with the repeated label as feedback, later, she expected the object to be displayed on the robot's tablet).

IV. ORIGINAL VALUE

Our analysis shows that PFs as recurring interaction patterns are also used in HRI. To benefit from PFs, we propose the approach of a pre-programmed base-PF on the level of dialog acts. The robot learner then would have to adapt this base-PF in interaction to represent its concrete realization. This could be realized with a calibration mechanism that starts with given knowledge and records the frame when this knowledge is taught.

REFERENCES

[1] Rohlfing KJ, et al., (2016) "An Alternative to Mapping a Word onto a Concept in Language Acquisition: Pragmatic Frames", Front. Psychol. 7:470, doi: 10.3389/fpsyg.2016.00470

[2] Vollmer A-L, et al. (2016) "Pragmatic Frames for Teaching and Learning in HumanRobot Interaction: Review and Challenges", Front. Neurorobot. 10:10, doi: 10.3389/fnbot.2016.00010

[3] SoftBank Robotics, https://www.softbankrobotics.com/emea/en/pepper (accessed June 19th, 2019)

[4] S. Garrido-Jurado, et al. (2014) "Automatic generation and detection of highly reliable fiducial markers under occlusion", Pattern Recogn. 47, 6, 2280-2292. DOI=10.1016/j.patcog.2014.01.005

[5] Bruner, J. S. (1983). Childs Talk: Learning to Use Language. New York, London: Norton and Co.