Eager to Learn vs. Quick to Complain? How a socially adaptive robot architecture performs with different robot personalities

Ana Tanevska*†, Francesco Rea*, Giulio Sandini*, Lola Cañamero‡, Alessandra Sciutti§
*Department of Robotics, Brain and Cognitive Science, Italian Institute of Technology (IIT), Genova, Italy
Email: {ana.tanevska, francesco.rea, giulio.sandini}@iit.it
† DIBRIS, University of Genova, Italy
‡ Embodied Emotion, Cognition and (Inter-)Action (EECAiA) Lab, School of Computer Science
University of Hertfordshire, Hatfield, United Kingdom
Email: l.canamero@herts.ac.uk
§Cognitive Architecture for Collaborative Technologies Unit, Italian Institute of Technology (IIT), Genova, Italy
Email: alessandra.sciutti@iit.it

***Abstract*—A social robot that is aware of our needs and continuously adapts its behaviour to them has the potential of creating a complex, personalized, human-like interaction of the kind we are used to have with our peers in our everyday lives. We are interested in exploring how would an adaptive architecture function and personalize to different users when given different initial values of its variables, i.e. when implementing the same adaptive framework with different robot personalities. Would an architecture that learns very quickly outperform a slower but steadier learning profile? To further explore this, we propose a cognitive architecture for the humanoid robot iCub supporting adaptability and we attempt to validate its functionality and test different robot profiles.**

***Index Terms*—Social robots and social learning, Human-human and human-robot interaction and communication, Architectures for Cognitive Development and Open-Ended Learning**

I. Introduction

Most people have a natural predisposition to interact in an adaptive manner with others, by instinctively changing our actions, tones and speech according to the perceived needs of our peers. This means that we are not only capable of registering the affective and cognitive state of other people, but over a prolonged period of interaction we also learn which behaviours are the most appropriate and well-suited for each one of them individually. Such universal trait that we share regardless of our different personalities is referred to as adaptability. Humans are always capable of adapting to the others although our personalities may influence the speed and efficacy of the adaptation.

Bringing this trait of adaptability to human-robot interaction (HRI) would provide user-personalized interaction ,a crucial element in many HRI scenarios - interactions with older adults, assistive robotics, child-robot interaction (CRI) with a learning focus, etc [1][2][3][4][5]. In our past works, we have explored adaptation in CRI in the context of switching between different game-based behaviours [6]. The architecture was affect-based [7], and the robot expressed three basic emotions (happy, sad, and a ”neutral” state) in a simple way. These emotions were affected by the level of engagement the child felt towards the current robot's behaviour. The robot aimed to keep the child entertained for longer by learning how the child reacted to the switch between different game modalities.

We have expanded the core concept of a robot's internal state guiding the adaptation, and advanced from the discrete emotional states and one-dimensional adaptation to a more robust framework. Starting from the work of Hiolle and Cañamero [8][9] on affective adaptability, we have modified our architecture to utilize as motivation the level of comfort of the robot, which is increasing when the robot is interacting with a person, and decreasing when it is left on its own.

In addition to exploring how would our adaptive architecture function with different users, we are interested in how different robot profiles would perform in the same adaptive framework. To address this, we designed a study exploring how that same architecture will behave and adapt if it is initialized with different sets of parameters, i.e. with different robotic personalities.

The robotic platform selected for our study was the humanoid robot iCub [10], and the scenario for testing the framework's functionalities was inspired by the typical interaction between a toddler and its caregiver, where the toddlers tend to seek the attention of their caretakers after being alone for a while, but as soon as their social need has been saturated they lose interest and turn their attention to something else [11]. The robot therefore acted as a young child, asking the caretaker's company or playing on its own and the human partners could establish and maintain the interaction by touching the robot, showing their face and smiling. This scenario was suitable to study some fundamental aspects of interaction (such as initiation and withdrawal) with a fully autonomous behaviour from the robot and very limited constraints to human activities, in a seemingly naturalistic context.

The rest of the paper is organized as follows: Sect. II presents the adaptive framework for our architecture and the validation studies on it, with Subsect. II-A presenting the architecture design and Subsect. II-B presenting an experimental testing and validation of the architecture. This is followed by Sect. III which presents the simulation study in which we tested the performance of our architecture with nine different robot personalities. Finally, in Sect. IV and V we present the findings from our studies and we touch on our plans for future work.

II. Methods

A. Robot Architecture

The personalized adaptive architecture for the robot was designed following the requirements for artificial and natural cognitive agents as defined by Vernon in [12]. Cognitive agents need to be able to perceive the environment, adapt to new circumstances in the environment, learn from their experience and ultimately act with the purpose of achieving some inner goals.

Starting from this basis, our framework for the iCub consisted of the following modules and their functionalities:

- Perception module, analysing two groups of stimuli:
  - Tactile stimuli - data processed from the skin sensors on iCub's arms and torso, carrying information about the size of the touched area (expressed in number of *taxels* - tactile elements) and the average pressure of the touch.
  - Visual stimuli - images coming from iCub's eye camera, detecting the presence of a face and for extracting the facial expression of the person.
- Action module, tasked with moving iCub's joint groups.
- Adaptation and motivation module, in charge of regulating iCub's social need and adapting its inner variables, motivated by maintaining an optimal level of comfort.

To further elaborate, the implementation of these functionalities in the software architecture comprised of:

a) Perception module: The perception module was realized using iCub's middleware libraries [10] for processing the data from the skin covers on its torso and arms, and the open-source library OpenFace [13] for extracting the facial features of the person. Fig. 1 shows a snapshot of the processing of the facial features and the taxels.

The data from the OpenFace library were analyzed for obtaining the most salient action units [14] from the detected facial features - lowering/raising eyebrows, crinkling of nose and cheeks, and smiling/frowning. These action units were weighted accordingly before being sent to the perception module, with the presence of a smiling person being rated as 1.0, a neutral or contemplating face as 0.75, a distant face as 0.5 and a frowning or disgusted face as 0.25. This was done by assuming that in an interaction a smile would bring higher social comfort than a neutral expression, similarly seeing a neutral face would still be more comforting than a distant presence or a displeased person.

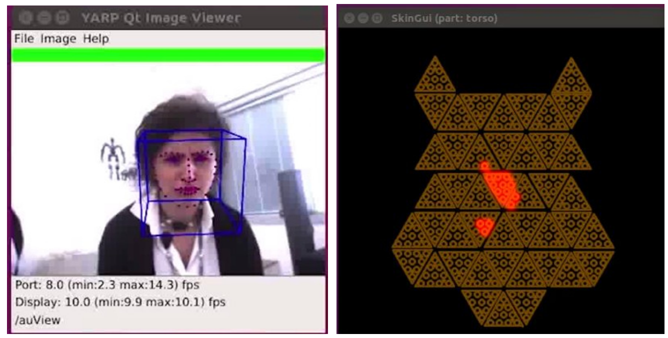

Fig. 1: Extracting facial features from iCub's cameras and tactile information from the skin on iCub's torso

The data from the skin sensors as well required some further processing post-extraction; due to the heating of the robot motors, phantom signals were registered during prolonged interaction. The output was filtered to register as touch only areas that were larger than 5 taxels and recorded average pressure larger than 12.0 (signifying that the average output of the activated taxels was registering a stable tactile contact). This data was processed for the torso and both arms separately, and sent to the perception module.

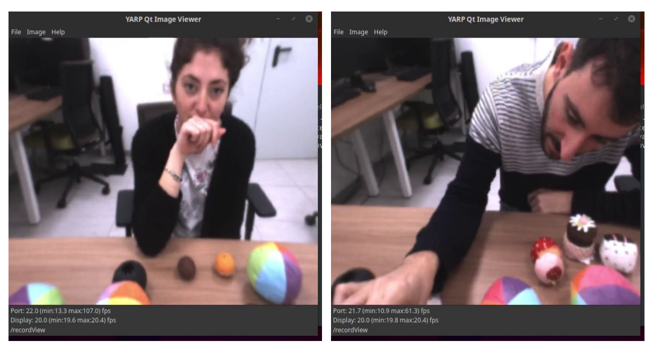

Fig. 2: Images from iCub's cameras during the validation study

b) Action Module: The action module communicated with iCub's middleware and performed a finite set of actions by controlling the specific body part in the joint space. These included head and neck motions determined by where iCub wanted to look, and arms and torso motions. If iCub wanted to engage with the caretaker, it would straighten up and look for the person, and then during the interaction engage in gaze-cueing and pointing to objects, whereas when iCub was oversaturated and wanted to disengage, it would pull away from the person and look down to its toys, ignoring other attempts to engage.

c) Adaptation and Motivation Module: This module maintained iCub's comfort and guided the adaptation process. The motivation in our architecture was represented by iCub's striving to remain in an optimal level of comfort. The comfort of iCub grew when a person was interacting with it, and the stimuli were weighted accordingly - a multimodal interaction (receiving both visual and tactile stimuli) or a longer, steadier interaction was rated higher and increased the comfort faster. Inversely, lack of any stimuli caused the comfort value to decay. iCub's social architecture was also equipped with a saturation and a critical threshold, which were reached when the interaction was getting too intense or was too sparse/non-existent, respectively.

At the beginning of the interaction with each user, iCub started with its comfort set at 50% of the maximum value it could have. Then the comfort level was updated continuously at the beginning of each cycle of the control loop of the interaction[1]. This happened in the following manner:

```
if (F[t] || T[t])
   C[t] = (F[t]+T[t]+C[t-1]*tau)/(tau+1)
else
   C[t] = beta*C[t-1]
```

If there was a person interacting with the robot (iCub was perceiving a face in front of it (`F[t]`), or registering touch with its skin (`T[t]`)), the comfort at the given moment in time `t` (`C[t]`) was updated as growing, taking into consideration both modalities in which the user could interact with iCub, as well as the of comfort in the previous instance of time (`C[t-1]`); inversely if iCub was not interacting with the user, the decay of the comfort was calculated. `tau` and `beta` were the growth and decay rates respectively. `tau` indicated how much the previous level of comfort was taken into consideration: a smaller `tau` meant a more rapid onset of the comfort when stimuli were detected, whereas a larger value meant a slower, steadier growth. `beta` was indicating how fast the comfort value decayed when there were no stimuli perceived; the smaller the value of `beta`, the more drastic the decay of the comfort. The initial values for the two rates were selected in consideration of the time it would take for an extreme user profile to reach the critical or saturation point, which was chosen as thirty seconds.

iCub's architecture allowed for adaptation on two dimensions - the frequency of interaction initiation and the duration of the interaction. The first one affected the decay rate of the comfort, and the adaptation on the second dimension instead modulated the growth rate of the comfort value. After each instance of iCub adapting on either dimension, it entered a suspension period where it attempted to recover and during which it was not open to interaction with the users.[2] The adaptation process had the following pattern:

- If the comfort reached the saturation limit: increase the value of `tau` (adapt with a slower comfort growth), and during the period of suspension ignore all stimuli. The resulting lack of sensitivity to stimulation leads to a decrease in the comfort value back to the optimal zone.
- If the comfort dropped to the critical level: increase the value of `beta` (adapt with a slower comfort decay), and during the suspension period simulate stimuli to itself so as to recover back to the optimal comfort level.

[1]Referring here to the perception-action control loop of iCub's architecture

[2]Originally the architecture adapted by immediately resetting the comfort level back to the optimal level and continuing with the interaction, which can be seen in Figure 3. The suspension period was included as a factor only after the validation of the original architecture with subjects, during which we realized that a continuation of responsiveness of the robot might not have allowed for the participants to infer that they were doing something not ideal for the robot. E.g. in the case of saturation, after the instantaneous robot withdrawal, it was immediately ready again to respond, which induced participants again to continue to interact in the same manner and trigger again saturation. The effect of different suspension periods is analysed in Sect. IV

B. Functional testing and validation

For the purpose of testing the architecture's functionalities, we defined three highly-varied user profiles - a highly-interactive profile that constantly attempted to engage in interaction with iCub, providing a very salient interaction that oversaturated the robot; a very sparsely interactive user who avoided mutual gaze and only engaged in tactile interaction once; and an in-between "mixed" profile that had periods of salient and sparse interaction, with the tactile one being more dispersed over the whole duration, and the mutual gaze and visual interaction happening only on one longer occasion in the last third of the interaction.

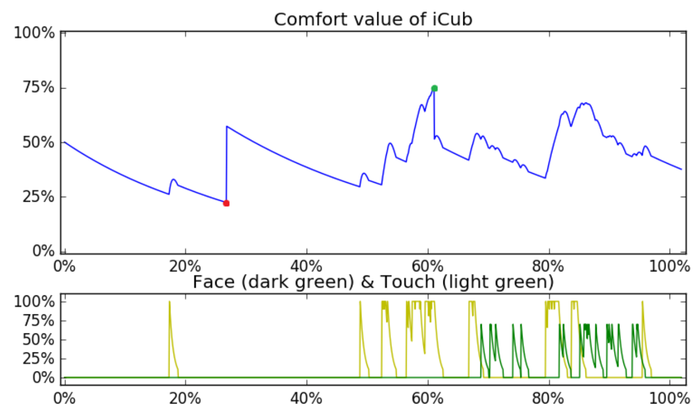

Fig. 3: Variations of the comfort value, with the x-axis showing the progress of the interaction in percentage

Fig. 3 shows the behaviour of the architecture for the profile with mixed interaction. The moments of received stimuli by the user are shown on the lower graph. The upper graph depicts the comfort level of iCub and how it decayed and grew during the interaction. In this profile there was one moment when the comfort was saturated (shown as a green dot on the peaks) and another instance when the iCub was left without stimuli for a while and it reached a critical value (shown as a red dot over the lowest point of the comfort). The adaptation of the decay rate can be seen as the decay gets slower after the critical trigger (the red dot), a similar observation is evident also for the growth rate before and after the saturation adaptation (the green dot). It is also evident how the strength of the different stimuli affect the comfort level, with it growing more rapidly when there are multiple inputs present at the same time. Different comfort trends were produced in response to the different simulated user profiles, with more pronounced growth rate changes in the highly interactive profile and more decay rate adaptations for the sparsely interactive one, however all showing the potentiality of the architecture to adapt and optimize the robot's behaviour.

After the first testing of the architecture, we additionally ran a small validation study [15] with naive users to assess whether the framework could successfully guide a real interaction. The participants (3 female and 3 male, mean age 28 +/- 4.89) interacted with iCub as its caretaker in two sessions, one where iCub had a fixed, scripted behaviour without an adaptive framework, and the other where iCub was behaving according to the adaptive personalized architecture described in Sect. II-A.

The results showed that the architecture could guide the interaction and make it pleasant for the human partner - in most cases more pleasant than when the robot behaviour was scripted. This study was very useful also to infer which are the natural ways in which participants behave toward a robot in a context like this one. Different individuals exhibited very different behaviours, which in turn affected very differently the evolution of the internal states of the robot derived from the architecture.

III. Simulation study

After testing and validating the functionalities of the architecture as well as trying it out in real-world interaction with participants, we approached our main research question - how would our architecture adapt to different users when given different initial values of its parameters, i.e. working with different robot profiles? If the considered parameter is the learning rate, would an adaptive robot who is a very fast and eager learner (i.e. takes big steps in the adaptation process) overshoot and miss the chance for personalization? If the considered parameter is the initial threshold value, would a finicky/fussy robot who has very narrow thresholds for interaction (i.e. has a very small difference between its saturation and critical thresholds) be an annoying interaction partner?

As we wanted to test multiple sets of parameters across multiple user profiles, we opted for designing five simulated user profiles and running a simulation study (before ultimately proceeding to full user studies), so that we would obtain the exact user behaviour across all conditions. For this we leveraged on the data obtained from the validation study in II-B to study how different users interacted with the robot [15].

TABLE I: Features of the five user profiles

	Visual	Tactile	Answer call
Complete (c)	1.0	1.0	always
Frequent (f)	[0.75-1.0],20 sec.	[0.0-1.0],15 sec.	always
Average (a)	[0.5-0.75],60 sec.	[0.0-1.0],10 sec.	always
Sparse (s)	[0.0-0.5],20 sec.	[0.0-1.0],15 sec.	only once
Void (v)	0.0	0.0	never

Table I showcases the modalities of the profiles. The visual stimuli were designed to be alternated between two values on a fixed time intervals, and the value in seconds shows at which frequency they alternated; the tactile stimuli instead were given depending on the state of the robot, the *average* and *sparse* profiles only provided tactile stimuli as a response to a call for engagement, whereas the *frequent* profile also provided stimuli while the robot was in an interactive state, and the value in seconds shows how long the tactile contact was. The *complete* and *void* profiles had either constant or non-existent input.

A. Implementing the suspension period

As we mentioned in II-A, the current version of the adaptive module leveraged on the functionality of a suspension period in which iCub could "recover" from the failed interaction. Before proceeding with the full simulation study, we ran a preliminary investigation across the user profiles to determine the optimal length of the suspension period. We selected as potential values the durations of 5, 20 and 35 seconds, and we tracked the effect the suspension period had on two metrics - the amount of time (expressed as a percentage from the whole duration of the interaction) in which iCub was in an optimal zone of comfort (i.e. the comfort value was within a maximum of 5% distance of the saturation and critical thresholds), and the number of times it reached a critical or saturation threshold and went into the adaptive module. In our framework we wanted to maximize the former and minimize the latter. The minimization of the number of adaptation signified a smaller amount of interruptions in the interaction flow by iCub disengaging, whereas maximizing the time iCub was in the optimal zone of comfort ensured an interaction where iCub would be neither too annoying to the person by constantly asking for attention, nor too isolated by not tolerating a sparsely interactive user.

TABLE II: Result per different suspension period lengths. Lighter grey highlights the best results per profile, and in darker grey is shown the optimal length.

		Suspend period (sec)		
metric	profile	5	20	35
Comf. %	c	47.54	83.24	71.94
	f	61.12	81.35	74.51
	a	60.13	69.97	59.18
	s	91.56	82.11	73.26
	v	40.55	68.32	65.87
Adapt #	c	8	4	5
	f	6	4	5
	a	7	7	7
	s	1	1	5
	v	8	5	6

Note: Comf % - percentage of total interaction time in which the robot is in the optimal comfort level. *Adapt #* - number of triggered adaptations during the interaction.

B. Testing different robot profiles

After running the suspension optimization study, we implemented the suspension interval that gave the best results in our adaptation module and proceeded to design the set of robot profiles for the simulation study, i.e. select the sets of parameters we would want to explore.

We experimented on two dimensions - the speed of adaptation, which was the step of modification for the growth and decay rate in the adaptation module; and the width of the band between the two thresholds. In terms of robot personalities, this could be translated to experimenting between a slow and

TABLE III: Results for different step sizes for adaptation and different thresholds distance (band width). The best results for each user profile are shown in light grey and the best results over all profiles are highlighted in dark grey

step size		slow			medium			fast		
band width		25%	50%	75%	25%	50%	75%	25%	50%	75%
Comf %	c	72.46	83.25	77.88	73.48	84.64	78.86	70.79	82.13	85.76
	f	71.06	81.96	80.08	71.70	82.34	84.86	69.65	79.81	82.27
	a	22.83	75.49	83.81	16.76	82.57	84.58	20.04	68.37	77.67
	s	28.19	81.83	85.34	39.14	81.44	84.90	51.75	81.14	77.21
	v	57.32	68.34	64.23	15.47	75.56	64.88	66.96	77.99	61.84
avg		50.37	78.17	78.27	43.31	81.31	79.62	55.84	77.89	76.95
Adapt #	c	5	4	4	4	3	3	4	3	2
	f	5	4	3	4	3	2	4	3	2
	a	5	5	7	2	3	6	2	3	6
	s	5	1	3	2	1	3	4	1	4
	v	3	5	8	2	3	7	1	2	7
avg		4.60	3.80	5.00	2.80	2.60	4.20	3.00	2.40	4.20
Idle %	c	3.06	3.18	2.96	3.19	3.29	3.45	3.16	3.33	3.47
	f	3.06	3.10	3.33	3.20	3.17	3.47	2.89	3.62	3.15
	a	33.26	30.82	33.10	61.30	39.60	38.59	45.81	55.99	48.02
	s	11.27	10.14	7.77	18.67	10.04	8.36	20.67	10.14	8.32
	v	44.00	38.20	24.06	21.09	40.63	26.92	41.96	39.49	28.82
avg		18.93	17.09	14.24	21.49	19.35	16.16	22.90	22.51	18.36
Int %	c	55.71	62.07	68.34	61.62	68.45	75.69	62.11	68.69	75.75
	f	55.99	62.69	68.57	61.80	69.17	76.31	62.81	68.43	76.42
	a	63.19	65.66	62.85	36.10	57.56	57.80	51.41	41.46	48.69
	s	48.46	73.45	58.09	47.67	73.61	58.00	40.61	73.60	56.17
	v	9.29	11.39	10.70	3.40	8.90	11.02	7.73	8.86	9.42
avg		46.53	55.05	53.71	42.12	55.54	55.76	44.93	52.21	53.29
Sus %	c	37.95	31.38	25.54	31.81	24.82	17.25	31.39	24.49	17.18
	f	37.66	31.15	24.61	31.62	24.50	16.94	31.24	24.47	17.18
	a	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	s	27.97	8.83	23.52	23.33	8.80	23.44	21.56	8.78	22.93
	v	29.03	32.77	43.12	9.82	25.20	42.83	25.09	25.00	42.70
avg		26.52	20.83	23.36	19.32	16.66	20.09	21.86	16.55	20.00

Note: Comf % - percentage of total interaction time in which the robot is in the optimal comfort level. *Adapt #* - number of triggered adaptations during the interaction. *Idle %, Int % and Suspend %* - percentage of total interaction time during which the robot is in one of these three states.

steady vs. a fast and eager learner (depending on the size of the adaptation step for the growth/decay rates); and between a very tolerant vs. a very fussy robot (depending on how wide the threshold band was, i.e. how close the saturation and critical points were to each other).

For both dimensions we chose three values to test, giving us a total of nine combinations of profiles. We explored three speeds of adaptation (slow, medium and fast) and three band widths (measuring in size as 25%, 50% or 75% of the total range of the comfort value of the robot, which could fall between 0.0 and 2.0). The length of the interaction sessions was set to ten minutes.

IV. Results

As it can be seen in Table II, the most optimal suspension period length was shown to be the one of 20 seconds, which nearly uniformly maximized the comfort of the robot (except for the sparse profile) and had the minimal amount of adaptations across all user profiles.

The simulation study was run on nine robot profiles and five simulated user profiles, giving us a total of 45 sessions, which were performed using a 20-second suspension interval, as derived from the previous simulation. The collected data was analysed for five metrics: the amount of time in which iCub was in an optimal zone of comfort (Comf %), the number of times iCub reached a critical or saturation threshold (Adapt #), and additionally the amounts of time during the interaction that iCub spent in an interactive, idle or suspended state (Idle %, Int % and Sus % respectively).

The results are shown in Table III. Unlike the simulation study for the suspension period, in this study there was not a single dominant "optimal" set of parameters. Depending on the metric we were interested in, there were different robot profiles which showed the best results.

On general, the slowly-learning robot did not outperform the other personalities in most of the metric categories, except for the % of time spent in an idle state. However, although with this robotic profile of a slowly-learning and highly-tolerant robot there was the least average amount of time spent in the Idle state, this was not due to increased interaction time, but an increased amount of adaptation hits. The robot being slowly learning meant that it hit a threshold limit more often since it adapted only by a small amount each time.

Overall, the best averaged results across profiles were found in the moderate or fast-learning profiles. The moderate learner had the best averaged result for the amount of time the robot was in its optimal comfort zone (81.31% of the interaction, medium step with 50% band width), as well as the second-smallest amount of adaptations, 2.6. The smallest average number of adaptations instead was for the fast-learning robot in the 50% band, 2.4.

The moderate learner also had the two highest % of Interaction times, 55.76% and 55.54% for the 50% and 75% bands respectively. Finally, having a very "fussy" robot with a very narrow threshold band did not produce optimal averaged results even at the fastest learning speed.

All robot profiles were initialized with parameters that made them call out to the user for contact if left alone (or withdraw from interaction if too stimulated) after approximately 30 seconds. After the experiment this changed in various directions as a function of the (simulated) user needs - the *complete* user profile influenced the adaptation exclusively in the saturation attitudes of the robot and slowing down considerably the growth rate, the *void* user profile affected only the decay rate, and the remaining profiles triggered adaptation events depending on their individual frequency and length of interaction. While at the beginning the *complete* and *void* profiles reached the first adaptation point in all robot profiles after approximately 30 seconds, at the end of the interaction these values had changed to range between 170-260 seconds for the *complete* profiles and 110-140 seconds for the *void* profiles. Apart from being a validation of the adaptation framework, these results also stress the additional benefit of having a two-dimensional adaptation and how it can contribute to the personalization of the robot's behaviour on more than one modality

V. Discussion and future work

Different individuals have different inclinations to interact with others and this applies also to their approach to interaction with robots. At the same time different tasks might require different level of human intervention (or robot request for help). Creating a unique robot behaviour (or personality) able to fit with task constraints and at the same time with individual desires is an impossible challenge. Endowing the robot with possibility to adapt to its partners' preferences is therefore important to grant a certain degree of compliance with individual inclinations. On the other hand, also the initial "personality" of the robot - i.e. the parameters at the beginning of the interaction - have a strong influence on the dynamics of the adaptation and on the appropriateness of robot resulting behaviour with respect to task constraints.

In this paper we have presented our personalized adaptive robot architecture, and the results from the simulation studies aimed at testing different sets of the architecture's parameters corresponding to the different robot profiles. We described how our architecture enables the robot to adjust its behaviour to suit different interaction profiles, using internal motivation which guides the robot to engage and disengage from interaction accordingly, while also taking in account the behaviour of the person interacting with it. We wanted to investigate how initializing the architecture with different values for its internal variables at the beginning (i.e. endowing the robot with different personalities) will affect the flow of the interaction and the extent of the adaptation. From the simulation study we noted that there was not one universally optimal set of parameters, i.e. robot profile that exceeded in performance across all evaluation metrics and all user profiles. However, at least in simulation, a medium-to-fast adaptation pace lead to a better interaction across almost all user profiles. Also, even with a very demanding robot (smaller band width) and slow paced adaptation it is possible to observe a change in robot behaviour over time, suggesting that even in case of not optimal selection of initial parameters, the adaptive process was able to progressively tune the interaction to the needs of the individual.

Continuing from the results of the simulation studies, we are interested in seeing how real users would respond to interactions with different robot personalities. We have already seen that when faced with an adaptive and static robot profile, users expressed preference for the adaptive robot [15]. A follow-up question will be to see how the different adaptive traits of the architecture would be received by the users, and whether the emergent behaviour we found in the simulation study will also carry to the one with users. Our hope is that our adaptive framework will provide for a more individualized, long-term, generalized interaction between humans and robots.

Acknowledgment

The authors would like to thank those who participated in the validation study, as well as Matthew Lewis and Imran Khan from the University of Hertfordshire for their help and availability. Results incorporated in this publication have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, G.A. No 804388, wHiSPER.

References

[1] Vaufreydaz D, Johal W, Combe C (2016) Starting engagement detection towards a companion robot using multimodal features. Robotics and Autonomous Systems 75:4–16

[2] Jimenez F, Yoshikawa T, Furuhashi T, Kanoh M (2015) An emotional expression model for educational-support robots. Journal of Artificial Intelligence and Soft Computing Research 5(1):51–57

[3] Breazeal C, Scassellati B (1999) How to build robots that make friends and influence people. In: Intelligent Robots and Systems, 1999. IROS'99. Proceedings. 1999 IEEE/RSJ International Conference on, IEEE, vol 2, pp 858–863

[4] Paiva A, Leite I, Ribeiro T (2014) Emotion modeling for social robots. The Oxford handbook of affective computing pp 296–308

[5] Cañamero L, Blanchard AJ, Nadel J (2006) Attachment bonds for human-like robots. International Journal of Humanoid Robotics 3(03):301–320

[6] Tanevska A (2016) Evaluation with emotions in a self-learning robot for interaction with children. Master's thesis, Faculty of Computer Science and Engineering, Skopje, Macedonia

[7] Tanevska A, Rea F, Sandini G, Sciutti A (2018) Designing an affective cognitive architecture for human-humanoid interaction. In: Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, ACM, pp 253–254

[8] Hiolle A, Cañamero L, Davila-Ross M, Bard KA (2012) Eliciting caregiving behavior in dyadic human-robot attachment-like interactions. ACM Transactions on Interactive Intelligent Systems (TiiS) 2(1):3

[9] Hiolle A, Lewis M, Cañamero L (2014) Arousal regulation and affective adaptation to human responsiveness by a robot that explores and learns a novel environment. Frontiers in neurorobotics 8:17

[10] Metta G, Sandini G, Vernon D, Natale L, Nori F (2008) The icub humanoid robot: an open platform for research in embodied cognition. In: Proceedings of the 8th workshop on performance metrics for intelligent systems, ACM, pp 50–56

[11] Feldman R (2003) Infant–mother and infant–father synchrony: The coregulation of positive arousal. Infant Mental Health Journal: Official Publication of The World Association for Infant Mental Health 24(1):1–23

[12] Vernon D (2014) Artificial cognitive systems: A primer. MIT Press

[13] Baltrušaitis T, Robinson P, Morency LP (2016) Openface: an open source facial behavior analysis toolkit. In: Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on, IEEE, pp 1–10

[14] Ekman P, Friesen WV, Hager JC (1978) Facial action coding system (facs). A technique for the measurement of facial action Consulting, Palo Alto 22

[15] Tanevska A, Rea F, Sandini G, Cañamero L, Sciutti A (in press) A cognitive architecture for socially adaptable robots. In: 2019 9th Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)