Long-term affect sensitive and socially interactive companions

Ginevra Castellano[1], Ruth Aylett[2], Kerstin Dautenhahn[3], Ana Paiva[4], Peter W. McOwan[1] and Steve Ho[2]

Abstract. This paper presents some of the requirements for the design of long-term affect sensitive and socially interactive companions. Research conducted in the EU project LIREC (LIving with Robots and intEractive Companions) envisage some capabilities including affect sensitivity, memory and learning, cognitive and expressive behaviour, personalisation and embodiment, highlighting the key issues that research on artificial companions should address.

1 INTRODUCTION

Machines endowed with emotional and social intelligence are becoming increasingly essential for many applications involving direct interaction with human users [1]. Artificial companions, whether as robots, graphical synthetic characters or socially interactive toys, are examples of artefacts that would benefit from the integration of social, affective capabilities into their underlying technology [2].

Artificial companions can be of crucial importance in many applications. As the average age of the population in many countries increases, health care for elderly people is becoming more problematic. Artificial companions will be able to provide additional functionalities that will assist primary and secondary users, e.g. carers, healthcare workers, etc. Companions could also represent valuable tools for edutainment and therapy applications, as well as impact the entertainment industry (e.g., design of socially intelligent toys, intelligent interactive games, etc.). They could act as personal assistants in smart environments and be employed as interactive toys for therapy and rehabilitation purposes, for example by encouraging and mediating interactions between people affected by social, cognitive disabilities (e.g., people with autism [3]).

So far, existing prototypes of artificial companions have had only limited functionality. While the possibilities opened up by digital technology are becoming larger and larger, computers and robots still lack many important social abilities and are not able yet to engage with humans in 'truly' natural interaction (as compared with human-human interaction). Establishing a relationship with human users requires an artificial companion that understands the way humans communicate, is able to infer their mental and affective states based on their verbal and non-verbal behaviour [4] and acts in an appropriate way. For example, a socially intelligent companion acting as a personal assistant would not bore a human user by trying to help her to accomplish a specific task if she is not in a good mood or is planning to engage in some other activity (thus a companion should not interrupt you watching your favourite TV programme). Artificial companions must be able to engage in long term interactions, to adapt to unforeseen circumstances over an extended period of time, and they must also be personalised, that is, know what the user likes or dislikes and what her preferences or her habits are. These are some of the most important characteristics that an artificial companion must have to be able to engage in a social interaction with human users.

This paper presents a brief review of key issues in this research field, specifically focusing on the research that is being carried out by the European project LIREC (LIving with Robots and intEractive Companions).

2 LONG-TERM AFFECT SENSITIVE AND SOCIALLY INTERACTIVE COMPANIONS

A robot companion can be defined as a robot that a) is useful, that is, able to help and assist people (though note the ability to entertain and motivate may also be an important task in many applications), and b) acts in a socially acceptable manner [5]. In the LIREC project we now take a further step and add the 'long-term perspective', i.e. we study companions that interact with people long-term and that may result in companion-human relationships. Note, the goal of LIREC is not to replace human contact, but to provide companions that fulfil their tasks and interact with people in a socially and emotionally acceptable manner. In scenarios where human-companion contact is long-term, we can expect that people will form relationships with such artificial companions. For this reason, LIREC specifically addresses the issue of long-term companionship from a multi-disciplinary perspective, in which the development of new technologies is accompanied by critical studies of psychological, ethological and ethical issues that arise from such developments.

Previous studies have shown that the novelty effect of artificial companions often quickly disappears [6] (see also [7] for an example of successful long-term companion). People tend to change their attitude and preferences towards the companions over time, and what they consider 'funny' or 'cool' initially may be perceived as 'boring' or 'annoying' at a later stage. Thus, artificial companions should be capable of adapting to the user's changes of attitude and state and behave accordingly in order to

[1] Department of Computer Science, Queen Mary, University of London, UK. Emails: {ginevra, pmco}@dcs.qmul.ac.uk.
[2] School of Maths and Computer Science, Heriot-Watt University, UK. Email: ruth@macs.hw.ac.uk
[3] Adaptive Systems Research Group, University of Hertfordshire, UK. Email: K.Dautenhahn@herts.ac.uk
[4] INESC-ID / Instituto Superior Técnico, Portugal. Email: Ana.Paiva@inesc-id.pt

keep up the user's interest level. Companion technology must also be designed, developed and tested to achieve long-term robustness in real world scenarios. The LIREC project aims to establish a multi-faceted framework for artificial long-term companions, embody it into innovative technology, evaluate the framework and technology in real environments and provide guidelines for designing such companions. In this paper we will focus on the following capabilities of companions: (1) affect sensitivity, (2) memory and learning, (3) cognitive and expressive behaviour, (4) personalisation, (5) embodiment.

2.1 Affect sensitivity

Affect sensitivity refers to the ability to analyse the verbal and non-verbal behaviour of users in order to understand their affective states.

Most previous studies focused primarily on the design of systems able to recognise basic emotions (e.g., joy, sadness, disgust, surprise, fear and anger), and were largely based on *acted* affective expressions [8]. While few studies have so far addressed the problem of finding methods for inferring more complex states, the design of artificial companions requires an affect sensitivity which goes beyond the ability to recognise prototypical emotions, and is able to capture spontaneous and more variegated affective signals conveying more subtle states such as, for example, boredom, interest, frustration, agreement, etc. [8]. Kapoor and colleagues, for example, designed a system that can perceive and respond in real-time to multimodal non-verbal cues that precede frustration in students using a learning companion [9].

Another important issue in relation to affect sensitivity is the need for multimodal systems that are able to fuse different channels of information in order to achieve a better understanding of the affective message communicated by the user. While unimodal systems (mainly based on facial expression or speech analysis) have been deeply investigated, studies taking into account the multimodal nature of the affective communication process are still not numerous. Some studies addressed bimodal affect recognition based on the fusion of facial expressions and head movement data (e.g., [10]), facial expression and body gestures (e.g., [11]), facial expressions and speech (e.g., [12]). A number of studies on multimodal affect recognition have also been reported in the literature. In [9] frustration is predicted using several affective expressions including facial expressions, head movement, posture, skin conductance and pressure data. Another study shows how facial expressions, body gesture and speech information is fused together at different levels to infer eight emotions [13].

A relevant issue in multimodal affect recognition is represented by the fusion of different modalities. Often features from different affective expressions are incompatible and their relationship is unknown. Several studies show how using combination schemes other than direct feature fusion allows for better performances to be achieved. Shan et al. [14], for example, proposed to fuse facial expressions and body gesture information at the feature level using Canonical Correlation Analysis, which captures underlying relationships between the feature sets in different modalities.

For an interactive companion to be able to interpret the user's state, the analysis of the dynamics of human behaviour is a factor of crucial importance. Affective expressions evolve over time, together with the meaning they convey, and their dynamic changes communicate more than a simple static affect display. Castellano et al. [15] found that the timing of expressive motion cues (namely the attack and release of the temporal profile of the velocity of the head and the quantity of motion of the upper body) is important in explaining emotional expression in piano performances. Valstar et al. [16] showed the importance of the dynamics of face, head and shoulder expressions in distinguishing spontaneous from posed smiles.

Finally, an affect recognition system for an interactive companion has to be designed so as to be robust in real-life applications. Face and body detectors and feature trackers need to be robust to real environment conditions like illumination changes, occlusions, dynamic backgrounds, etc.

2.2 Memory and learning

An artificial long-term companion must be endowed with a dynamic, adaptive framework for memory and learning. A key problem for a long-term companion is that remembering 'everything' would produce far too much data both in relation to storage capabilities, but more important, in relation to retrieval of relevant information when new interactions take place. The companion needs to know what to remember, but also what to forget based on information about user inputs and context.

A model linking long-term memory and working memory, with specific reference to emotional and autobiographical memory, is required for achieving long-term effectiveness. Such a model should include processes allowing external stimuli and internal emotional state of the companion to cause *reminding*, i.e., to retrieve memories which are relevant to the current situation.

Episodic memory, which organises temporal sequences of events, has been researched in studies on both robots and virtual agents. Such a memory, for example, has been used in relation to robot localisation and map-building to reduce state-estimation computations [17]. Ogino et al. [18] implemented a long-term episodic memory for a virtual robot; emotions were associated with memory elements which were used to support Human-Robot Interaction (HRI) in a simple game.

The concept of an 'interaction history' [19] [20] involves the "temporally extended, dynamically constructed and reconstructed, individual sensorimotor history of an agent situated and acting in its environment including the social environment". However, this raises issues of representation and level: to date, in most of the work in the domain of robotics, memory is represented in a form that is close to sensors and actuators rather than symbolically, so as to make it easy to apply to sensori-motor coordination. This raises problems where the task is communication with human users at the symbolic level.

Work in modelling a complete human episodic memory (e.g., [21] [22]), establishes a common structure that consists of context, contents and outcomes/evaluation for agents to remember past experiences. These models were created to focus on the following three different aspects:

- Accuracy – how relevant situations can be retrieved from the memory
- Scalability – how to accommodate a large number of episodes while not decreasing significantly the performance of the system
- Efficiency – how to optimise the storage and recall of memory contents

Brom et al. [23] attempted to create a full episodic memory storing more or less everything happening around a virtual agent for the purpose of storytelling. This was used to answer specific questions about the agent's personal history by human users in real time. Forgetting processes were also partially implemented - in the agent's long-term memory records, less emotionally interesting records were deleted.
LIREC will draw on the previous research of team members [24] aimed at modelling the psychological concept of autobiographic memory computationally. This has already been integrated into a synthetic agent architecture. With this memory included, agents are not only capable of recognising and ranking significant events originating in their own experiences, but can also remember, recall and learn from these experiences. This both increases agent believability and produces a richer interactive experience for the user.

2.3 Cognitive and expressive behaviour

An artificial companion must have mechanisms for verbal and non-verbal social interaction. These mechanisms should include communicative expressions that can be used for the generation of expressive behaviour of the companion and must integrate non-verbal behaviour (e.g., facial and bodily expression of the companion) with speech capabilities. The expression of such a companion may impact not only on its believability but also influence how well the companion is accepted by the user over a long period of time. Thus, fine-tuned expressions (both facial and body, depending on the embodiment of the companion) and speech need to be investigated and created focusing on long-term interaction. Yet, the design of a companion should also consider the integration of this expressive behaviour with an architecture which includes memory, emotion, personality, adaptation, and autonomous action-selection. Previous work on agent architectures for characters, such as FAtiMA [25] or EMA [26], aimed at generating believable social behaviour addressing issues such as emotion modelling and its link with behaviour generation. Yet, the sustainability of the interactions with users require that we develop adjustable models and mechanisms that support both collaborative and autonomous decision-making, influenced by the companion's internal state and, most importantly, past experiences. More concretely, we need to incorporate the following components and use them to impact decision-making:

- A model of the user's personality, goals, beliefs, emotional state
- Social Relations – an explicit model of the relations between the companion and the user or other agents.
- Emotions – emotions experienced by the companion
- Personality – a model of the companion's personality
- Autobiographical Memory – memory of past personal experiences

Such an emotionally-sensitive framework, together with the companion's affect sensitivity, is the main requirement for a companion to be able to act empathically towards the user.

2.4 Personalisation

Personalisation refers to the ability, for a companion, to adapt to a specific user over time. This requires the companion to be endowed with a model of the user's personality, goals, beliefs, emotional state, which can be partially generated a priori and then dynamically and automatically adapted. An artificial companion must know the human user it is interacting with, what the user likes and dislikes (e.g., proxemic preferences), what her current affective or mental state is, in order to be able to generate a personalised response. Recent research in HRI on individual user differences has focused primarily on the role of individual differences as regards proxemic preferences and robot navigation in the presence of humans. Moreover, previous research has shown that analysis of proxemic data from robot sensors is capable of clustering different users according to personality type [27]. Yet, for long-term interactions, personalisation will have to go beyond the characterisation of the user by the personality type, but also infer and analyse specific preferences, ways of acting and interaction styles of the user, in order to really convey the adequate personalisation, so important for a companion.

2.5 Embodiment

Developing a framework of socially-acceptable embodiment is a topic of current interest in the research on artificial companions. Different types of embodiment can have a different impact on the user experience and affect the user's level of engagement over time. Concerning robot appearance, there is evidence, e.g. from previous studies conducted within the FP6 European project Cogniron (The Cognitive Robot Companion), that most people tend to prefer humanoid robots, but that individual differences are significant enough not to be able to conclude that a robot with human-like appearance and attributes necessarily make the interaction with a robot more attractive [28]. This seems to suggest that the personalisation aspect should be taken into account into the design of embodiment of artificial companions.
A particular study investigated the preferences of seventy-nine undergraduate participants for robot appearance presenting three different types of robot appearances: a basic appearance, a mechanoid appearance and a humanoid appearance [28] [29]. Moreover, the static appearance (ratings based on photographs of the robots) was compared with the dynamic appearance (ratings based on 'watching the robot in action', i.e., after exposure to video clips showing the functioning of the three different robot appearances in the same 'attention seeking' home scenario). In the videos, the different robots differed in overall appearance, in particular in the number of human-like features, as well as in their attention-seeking behaviour from the human (using voices or beeps and different arm gestures). Specifically, the humanoid robot had a human-like arm, a human voice and a (compared to the other appearances) detailed head with a number of human-like features. The mechanoid robot had a simple arm, a mechanical voice and a simple head. The basic robot had a simple gripper, a head consisting of a camera only, and used beeps in order to attract attention. Note, this scenario involved little direct HRI, and so it was decided to carry out video trials in contrast to live trials. Comparative studies have demonstrated that, at least in scenarios that do not involve extensive human-robot interactive and communicative engagement, similar results can be obtained with video-based and live trials [30] [31].
In order to provide an ecologically appropriate context the videos were shot in the University of Hertfordshire Robot

House, an environment furnished like a domestic environment. Feedback from participants in live HRI studies have shown that they feel much more at ease and relaxed compared to studies carried out on University premises. Questionnaire results from the three different robot appearances show that the basic appearance was neither particularly liked nor disliked, while the mechanoid appearance was mildly disliked. Overall participants were more favourable of the humanoid appearance and behaviour. This increase of ratings from mechanoid to basic to humanoid appearance was found to be consistent for both dynamic and static appearance ratings, which can be related to the left-hand side of Mori's uncanny valley hypothesis that suggests that, initially, an increase in human-likeness will result in an increase in familiarity of robot appearance (until a point is reached where too much human-likeness, combined with still noticeable non-humanness, results in an uncanny impression). See a detailed discussion in [32] on the uncanny valley.

The issue of robot appearance and behaviour is closely related to humans' perception of robots in terms of robot personality and the robot's tasks and context of use. For example, a survey by Khan [33] reported on people's preferences for a mechanical looking robot as a service or assistant robot at home. Also, a more serious task may result in preferences for a robot that presents itself similarly as more 'serious' [34]. The perception of the robot also influences how people respond to and interact with robots [35].

In summary, several studies have shown that a) robot appearance matters in how people perceive robots and interact with them, and b) robot appearance and its perceived personality will benefit from being matched to the robot's tasks and role in HRI studies. See Walters [29] for an in-depth discussion of these issues.

Note, an important issue in this discussion concerns the personality of the participants in the HRI studies. In the above mentioned study involving the basic, mechanoid and humanoid robot appearances, analysis showed that participants with low Emotional Stability and Extraversion scores tend to show preferences for the mechanoid robot appearance [36]. Generally, an impact of participants' personality traits on HRI trials has also been confirmed by other research groups and is currently an actively researched topic [37].

To conclude after this brief discussion of robot appearance, robot personality and participants' personality, it is at present difficult to provide detailed recommendations on how robots should look and behave in certain scenarios, in particular scenarios that have not been studied yet experimentally. The literature provides some starting points, but within the project LIREC all these issues, and others concerning embodiment that could not be covered in this brief section, will be explored in depth.

3 CHALLENGES AND CONCLUSION

This paper provides an overview of some of the capabilities that a long-term affect sensitive artificial companion should be endowed with. Beside the specific challenges behind the design of each capability, it stands to reason that a common issue is represented by the need to create technology which is robust in real world scenarios over a long period of time. Dynamic models which allow for the companions to dynamically perceive external inputs and process, interpret and use them to update their internal state and models and to plan an appropriate reaction are also needed. Dynamics is particularly important, since for an artificial companion to be able to work long-term it is necessary to continuously adapt its mind.

4 ACKNOWLEDGEMENTS

This work was supported by the EU FP7 ICT-215554 project LIREC.

REFERENCES

[1] R. W. Picard. Affective Computing. MIT Press (1997).

[2] C. Breazeal. Emotion and sociable humanoid robots. E. Hudlicka (eds.), *International Journal of Human-Computer Studies*, 59(1-2), pp.119-155 (2003).

[3] B. Robins, P. Dickerson, P. Stribling and K. Dautenhahn. Robot-mediated joint attention in children with autism: A case study in robot-human interaction. *Interaction Studies* 5:2, pp. 161-198 (2004).

[4] M. Pantic, N. Sebe, J. F. Cohn, and T. Huang. Affective Multimodal Human-Computer Interaction. *In Proc. ACM International Conference on Multimedia*, pp. 669-676, Singapore (2005).

[5] K. Dautenhahn. Socially Intelligent Robots: Dimensions of Human-Robot Interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480), pp. 679-704 (2007).

[6] Z.-J. You, C.-Y. Shen, C.-W. Chang, B.-J. Liu and G.-D. Chen, A Robot as a Teaching Assistant in an English Class. *In Proc. of the Sixth Int'l Conf. on Advanced Learning Technologies* (2006).

[7] F. Tanaka, A. Cicourel, J. R. Movellan. Socialization between toddlers and robots at an early childhood education center. PNAS 104 (46), pp. 17954-17958 (2007).

[8] Z. Zeng, M. Pantic, G.I. Roisman and T.S. Huang. A Survey of Affect Recognition Methods: Audio, Visual, and Spontaneous Expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, TPAMI-2007-09-0496, To Appear.

[9] A. Kapoor, W. Burleson and R. W. Picard. Automatic Prediction of Frustration. *International Journal of Human-Computer Studies*, 65(8), pp. 724-736 (2007).

[10] R. El Kaliouby and P. Robinson. Generalization of a Vision-Based Computational Model of Mind-Reading. *In Proc. 1st International Conference on Affective Computing and Intelligent Interaction*, Beijing, China (2005).

[11] H. Gunes and M. Piccardi. Bi-Modal Emotion Recognition from Expressive Face and Body Gestures. *Journal of Network and Computer Applications*, 30(4), pp. 1334-1345, Elsevier (2007).

[12] C. Busso, Z. Deng, S. Yildirim, M. Bulut, C.M. Lee, A. Kazemzadeh, S. Lee, U. Neumann and S. Narayanan. Analysis of emotion recognition using facial expressions, speech and multimodal information. *In Proc. of International Conference of Multimodal Interfaces*, pp. 205-211 (2004).

[13] G. Castellano, L. Kessous and G. Caridakis. Emotion recognition through multiple modalities: face, body gesture, speech. In: Peter C., Beale R. (eds.): *Affect and Emotion in Human-Computer Interaction*. LNCS, vol. 4868. Springer, Heidelberg (2008).

[14] C. Shan, S. Gong, and P. W. McOwan. Beyond Facial Expressions: Learning Human Emotion from Body Gestures. *In Proc. Of British Machine Vision Conference (BMVC'07)*, Warwick, UK (2007).

[15] G. Castellano, M. Mortillaro, A. Camurri, G. Volpe, K. Scherer. Automated analysis of body movement in emotionally expressive piano performances. *Music Perception*, University of California Press. In press.

[16] M.F. Valstar, H. Gunes and M. Pantic. How to Distinguish Posed from Spontaneous Smiles using Geometric Features. *In Proc. of ACM Int'l Conf. Multimodal Interfaces (ICMI'07)*, pp. 38-45, Nagoya, Japan (2007).

[17] Y. Endo. Anticipatory robot control for a partially observable environment using episodic memories. Technical Report GIT-IC-07-03, Georgia Tech Mobile Robot Lab (2007).

[18] M. Ogino, T. Ooide, A. Watanabe, and M. Asada. Acquiring peekaboo communication: Early communication model based on reward prediction. In *Proc. of IEEE International Conference in Development and Learning (ICDL)*, London, UK (2007).
[19] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn and R. te Boekhorst. Interaction histories: From experience to action and back again. In *Proc. of the 5th IEEE International Conference on Development and Learning (ICDL 2006)*, Bloomington, IN (2006).
[20] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn and R. te Boekhorst. Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny. *Adaptive Behavior 15(2)*, pp. 167–187 (2007).
[21] A. Nuxoll and J. E. Laird. A cognitive model of episodic memory integrated with a general cognitive architecture. In *Proc. of the Sixth International Conference on Cognitive Modeling*, Mahwah, NJ, Lawrence Earlbaum, pp. 220–225 (2004).
[22] D. Tecuci and B. Porter. A generic memory module for events. In *Proc. of the 20th Florida Artificial Intelligence Research Society Conference (FLAIRS 20)*, Key West, FL (2007).
[23] C. Brom, K. Peskova and J. Lukavskyz. What does your actor remember? Towards characters with a full episodic memory. In Cavazza, M., Donikian, S. (eds.): *International Conference on Virtual Storytelling.* Lecture Notes in Computer Science 4871, Springer, pp. 89–10 (2007).
[24] W. C. Ho, J. Dias, R. Figueiredo and A. Paiva. Agents that remember can tell stories: integrating autobiographic memory into emotional agents. In *Proc. of Autonomous Agents and Multiagent Systems (AAMAS)*, ACM Press (2007).
[25] J. Dias and A. Paiva. Feeling and Reasoning: A computational Model for Emotional Characters. In Bento, C., Cardoso, A. And Dias, G. (eds.): *Progress in Artificial Intelligence, EPIA'2005*, Springer, LNAI 3808 (2005).
[26] J. Gratch and S. Marsella. A Domain-independent Framework for Modelling Emotion. *Journal of Cognitive Systems Research*, 5(4) (2004).
[27] T. Salter, K. Dautenhahn and R. te Boekhorst. Learning about natural human-robot interaction styles. *Robotics and Autonomous Systems* 54(2), pp. 127-134 (2006).
[28] M.L. Walters, D. S. Syrdal, K. Dautenhahn, R. te Boekhorst and K. L. Koay. Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots*, 24(2), pp. 159-178 (2008).
[29] M. L. Walters. The design space of robot appearance and behaviour for social robot companions. PhD thesis, University of Hertfordshire, February 2008.
[30] S. N. Woods, M. L. Walters, K. L. Koay and K. Dautenhahn. Comparing Human Robot Interaction Scenarios Using Live and Video Based Methods: Towards a Novel Methodological Approach, *In Proc. AMC'06, The 9th International Workshop on Advanced Motion Control*, March 27-29, Istanbul (2006).
[31] S. N. Woods, M. L. Walters, K. L. Koay and K. Dautenhahn. Methodological Issues in HRI: A Comparison of Live and Video-Based Methods in Robot to Human Approach Direction Trials. *Proc. The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN06)*, University of Hertfordshire, 6-8 September, Hatfield, UK, pp. 51-58, IEEE Press (2006).
[32] K. MacDorman and H. Ishiguro. The Uncanny Advantage of Using Androids in Coognitive and Social Science Research. *Interaction Studies*, 7(3). pp. 297-337 (2006).
[33] Z. Khan. Attitude towards intelligent service robots. Technical report, TRITA-NAP9821, NADA, KTH (1998).
[34] J. Goetz and S. Kiesler. Cooperation with a Robotic Assistant. *In Proc. of the Conference on Human Factors in Computing Systems (CHI'02)*,New York, USA, pp. 578 -579 (2002).
[35] P. J. Hinds, T. L. Roberts and H. Jones. Whose Job Is It Anyway? A Study of Human-Robot Interaction in a Collaborative Task. *Human Computer Interaction*, 19. pp. 151-181 (2004).
[36] D. S. Syrdal, M. L. Walters, K. L. Koay, S. N. Woods and K. Dautenhahn. Looking Good? Appearance Preferences and Robot Personality Inferences at Zero Acquaintance. *AAAI – Spring Symposium 2007, Multidisciplinary Collaboration for Socially Assistive Robotics* (2007).
[37] A.Tapus, C. Tapus and M. J Mataric'. User-Robot Personality Matching and Robot Behavior Adaptation for Post-Stroke Rehabilitation Therapy, *Intelligent Service Robotics Journal, Special Issue on Multidisciplinary Collaboration for Socially Assistive Robotics*, 1(2):169-183 (2008).