

Lifelong robot-assisted mobility: Models, Tools, and Challenges

Yiannis Demiris and Tom Carlson

Assistive Robotics Laboratory, Department of Electrical and Electronic Engineering, Imperial College London
y.demiris@imperial.ac.uk, tom.carlson02@imperial.ac.uk
<http://www.iis.ee.ic.ac.uk/yiannis>

Keywords: Assistive Robotics, Lifelong User Models, Adaptive Shared Control, Intention Prediction.

Abstract

Increasing the autonomy of users with disabilities through robot-assisted mobility has the potential of facilitating their sensorimotor and social development, as well as reducing the burden of caring for such populations in both inpatient and outpatient settings. While techniques for task-specific assistance exist, they are largely focused on satisfying short-term goals, utilising stationary user models. For lifelong users and particularly for those with rapidly changing sensorimotor skills (for example very young children), adaptive models that take into consideration these developmental trajectories are becoming very important. In this paper, we present our approach to lifelong user models for robot-assisted mobility, and discuss existing models and tools, as well as challenges that remain ahead.

1 Introduction

Sensorimotor skills develop rapidly, especially during the first five years in children, and following a stroke or accident in adults. These skills are crucial for enabling developing children and adults under rehabilitation to explore their environment and engage in meaningful social interaction. Disabled adults as well as children with sensorimotor disabilities (for example due to cerebral palsy [6], spinal muscular atrophy [7], or tetraplegia [9] among others) lose crucial opportunities to use and develop their cognitive and social skills, which has a negative impact on their overall continued development and quality of life. Early interventions in the sensorimotor development of such populations through the use of powered mobility and computational mechanisms for adaptive shared control can facilitate socio-cognitive development and help approximate the development pattern of the non-disabled population. Increasing the independence of users through autonomous, safe powered mobility has also the potential to reduce the burden involved in caring for populations of disabled people in (frequently understaffed) clinical settings, such as rehabilitation and residential care centres. The introduction of powered mobility to such populations however is extremely challenging. For example, in the case of disabled children, despite documented benefits for the use of powered mobility (with earlier intervention preferable), therapists do not often consider it as an option,

citing post-provision difficulties with their use, and problems with the children learning initial control of the wheelchair, as well as a need to adjust operational parameters (such as speed and sensitivity of control hardware) according to the circumstances. In a national survey [5], 26 of 47 wheelchair service providers in England that had provided powered wheelchairs to children under the age of 5 reported post-provision difficulties, with 19 services specifically citing problems with the child learning initial control. On the socioeconomic side, time costs for parents caring for children with disabilities are significantly greater than those for non-disabled children, they do not decrease with age (but do relate with degree of disability), and have significant impact to the ability of parents to work outside their home [3]. For adult patients following stroke, positive changes in quality of life have been recorded following the provision of powered wheelchairs [11], but safety concerns frequently exclude potentially useful mobility solutions based on powered wheelchairs for certain groups of patients, for example those with unilateral neglect [10].

2 From user modelling to lifelong adaptive models

In our work we investigate how adaptive shared control can be utilised to enable users with disabilities (e.g. [1,2] for wheelchair-bound populations (figure 1)) to continue to safely operate their powered devices, and receive support only when needed, through a combination of prediction of human intention [4] and environmental situation awareness. In the following few subsections we review our experimental findings as a way for motivating and explaining the core challenges we have identified.

2.1 Obtaining the ‘user vector’

In the forefront of technical challenges towards intelligent assistive devices is the inevitable stage of extracting a collection of human data that will be used to calculate the current state of the human user of the device. While many human wheelchair users operate their devices through a joystick, this is not always adequate or even appropriate. In a survey of practicing clinicians [8], it was reported that 40% of patients find it difficult or impossible to engage in steering and manoeuvring tasks. The same survey reported that nearly

half of patients unable to control a power wheelchair by conventional methods would benefit from an automated navigation system. Alternative human machine interfaces that have been used for this task include the measurement of a number of physiological signals (such as EEG, EMG, EOG signals), and eye gaze (e.g. figure 2) or head movements.



Figure 1: An example setup at Imperial's assistive robotics laboratory; human desired commands are sent through the joystick to the shared control system on a tablet PC, which estimates (through the wheelchair's cameras and laser scanners) whether the resulting action will have the intended effect, and be safe for the user, and alters the commands appropriately.



Figure 2: A portable eye-gaze tracker mounted on bicycle helmet monitors the user's attention patterns helping us to quantify the effects of the collaborative control system on the user's perception [12].

2.2 Intention prediction

We use a generative approach to recognising the intentions of the human operator, utilising HAMMER (Hierarchical Attentive Multiple Models for Recognition and Execution) as the basis architecture [4]. HAMMER uses the concepts of inverse and forward models. An inverse model is akin to the concepts of a controller, behavior, action, or motor plan. The inverse model's function is to receive as input a measurement or estimate of the current state of the system and the desired target goal(s) and output the control commands that are needed to achieve or maintain those goal(s). A forward model of a system (akin to the concept of an internal predictor) is a function that takes as inputs the current state of the system and a control command to be applied to it and outputs the predicted next state of the controlled system. When HAMMER is asked to rehearse or execute a certain action to control the assistive robot, the corresponding inverse model module is given information about the current state and, optionally, about the target goal(s). The inverse model then outputs the motor commands that are necessary to achieve or maintain these implicit or explicit target goal(s). The forward model provides an estimate of the upcoming states should these motor commands be executed. The estimate can be compared with the target goal to produce a reinforcement signal for the inverse model depending on how much the model's motor commands brought the estimate closer to the target goal. The HAMMER architecture uses an inverse-forward model coupling in a dual role: either for executing an action, or for perceiving the same action when performed by a demonstrator/operator. When HAMMER operates in action perception mode, it can determine whether an ongoing action performed by a human user matches a particular inverse-forward model coupling by feeding the user's current state as perceived by the observer system to the inverse model. The inverse model generates the motor commands that it would output *if it was in that state and was executing the particular action*. HAMMER consists of multiple pairs of inverse and forward models that operate in parallel. As the human user executes a particular action, and there are multiple models (possibilities) that can explain the ongoing demonstration, we feed the perceived states into all of the system's available inverse models. This will result into the generation of multiple motor commands (representing the multiple hypotheses as to what action is being executed) that are sent to the forward models. The forward models generate predictions about the system's next state as described earlier and these are compared with the actual system's next state at the next time step. The error signals resulting from this comparison affect the confidence values of the inverse models. At the end of the action (or usually earlier as required in shared control (next)) the inverse model with the highest confidence value, i.e. the one that is the closest match to the operator's action is selected and is offered as an estimate of the intention.

2.3 Adaptive Shared Control

Using this intention prediction, the amount of shared control that will be given is calculated by activating the corresponding inverse model, deriving the series of motor commands that are required to achieve this intention, and monitoring and enhancing the user commands to ensure successful (and primarily safe) completion. For example, figure 3 demonstrates our approach to adaptive shared control of robotic wheelchairs [1,2]. The localisation module estimates the wheelchair's current location and orientation, as well as the user's potential targets for movement (e.g. doors), while the joystick commands (or EEG/EMG signals) are used to derive predicted intentions, and the confidence we have in these predictions. Ideal (safe) trajectories are generated, and the user's input is modified by the shared controller before being sent to the robot's motor units, if necessary.




Figure 3: an example adaptive shared-controller for wheelchair-bound populations [1,2]

3 Challenges for lifelong assisted mobility

Assisting in the way we described in the previous section has an immediate effect to the performance of the individual operating the wheelchair, but it is unclear what the effects of this assistance will be in the long term.

3.1 Balancing immediate and long term needs

Our target approach does not aim to engineer autonomous systems that will *replace* the user's control, nor *constantly and unconditionally* assist the user. We want to develop lifelong models that attempt to *balance* the current needs of the user with the challenges that the user is capable of overcoming *with assistance from the robotic system*. Principled methods to approach this task are in their infancy, and require progress in fundamental theoretical and engineering issues. This includes finding ways to *continuously* measure the situated sensorimotor activity of users, using body sensor networks, and to approximate their current sensorimotor capabilities and *potential for achieving*

their intentions with and without help (akin to the concept of the Zone of Proximal Development (ZPD) by L. Vygotsky).

3.2 Approximating the user's Zone of Proximal Development

The approach our work follows uses HAMMER's distributed networks of inverse and forward models [4] as a starting point for these approximations, as these have proven to be very useful in the fields of motor control, action recognition and intention interpretation. Our systems attempt to form and customise hierarchical combinations of primitive inverse models, and compare them against the user's performance to determine its ZPD. The next few sections describe the theoretical aspects of this algorithm, and the challenges that need to be overcome.

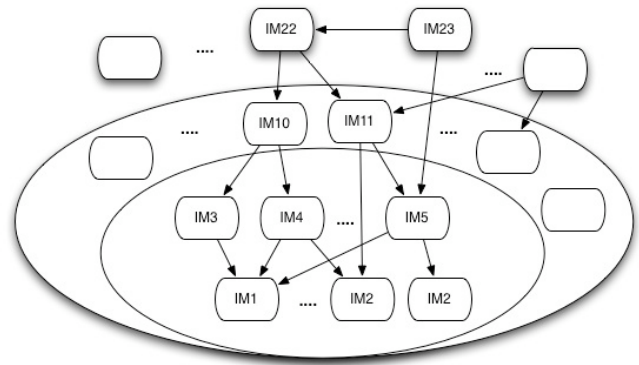


Figure 4: The zone of proximal development approximation using hierarchical inverse models; in this conceptual diagram, inverse models IM10 and IM11 are within the ZPD since they rely on IMs 3, 4, 2, and 5 which have been successfully executed previously.

3.2.1 Quantifying current performance: annotating IMs

Determining what the user is currently capable of is a very difficult task on its own; in our approach the amount of assistance given by the shared controller supplementing the user's commands is a good approximation of these capabilities, and the inverse models are annotated to indicate quantities and frequency of assistance given, as well as the effect it had on the overall user's performance (for example the resulting smoothness of the user's joystick movements [1,2]).

3.2.2 Approximating the potential for achieving the intention

Following the prediction of intention step, the HAMMER architecture has derived one (or more) inverse models (Intended Task's Inverse Model – ITIM) of what is required for the user to achieve his intention. An approximation as to whether this falls within his ZPD is now required; the ZPD approximation algorithm (ZPD-AA1) is as follows:

Algorithm ZPD-AAI

Input: The Intended Task's Inverse Model (ITIM)

Output: An estimation of whether the inverse model falls within the user's Zone of Proximal Development

1. Decompose ITIM into the component inverse models that it requires for its completion.
2. Retrieve the associated values for quantities and frequency of assistance given for the component inverse models in the past.
3. Approximate the predicted level of combined shared control that needs to be given. For this, either use heuristics (for example, if a number of component inverse models are organised serially, take the highest quantity of assistance among the components as approximation, while if there are components organised in parallel (need to be executed concurrently), use a weighted sum of the components as approximation), or learn this by observation.
4. Determine whether this falls within acceptable levels to qualify it as within the ZPD of the user. The threshold can become user- or practitioner-defined.

Having determined whether the intended task falls within the user's ZPD, the challenge that remains is to adapt the human-robot interface options and levels of shared control to facilitate execution and to maximise not only short-term benefits (success in the current task) but also long term development (such as deterioration of morale due to lack of opportunities to overcome challenges). Options include:

- A user-defined level of assistance that could be given in general (which the user can easily adjust to suit his current desires)
- A situation-dependent level of assistance, that can factor in the user's emotional state, contextual aspects (such as whether lack of assistance will pose a danger to self or others, or whether it will affect the perceptions of others towards the user (for e.g. assistance might be increased in public places and reduced in private settings).
- Disability specific training needs, potentially taking into consideration observations as to whether it was effective.
- Approaches that balance the last two options, based on practitioners' defined schedules.

4 Conclusions

Despite the difficulties in determining appropriate levels of shared control, algorithms that estimate and incorporate lifelong patient developmental trajectories have the potential to improve the benefits that assistive robotics can offer.

Acknowledgements

Thanks to Murilo Fernandes Martins, Mireno Rossi and Valerio Chang, for their help with the robotic experiments.

References

- [1] T. Carlson and Y. Demiris, "Human-Wheelchair Collaboration through Prediction of Intention and Adaptive Assistance", Proceedings of IEEE International Conference on Robotics & Automation, CA, USA, pp 3926-3931 (2008).
- [2] T. Carlson and Y. Demiris, "Collaborative Control in Human Wheelchair Interaction Reduces the Need for Dexterity in Precise Manoeuvres", In Proc. of ACM/IEEE HRI-2008 workshop on "Robotic Helpers: User Interaction, Interfaces and Companions in Assistive and Therapy Robotics", Amsterdam, Netherlands, pp 59-66. (2008)
- [3] A. L. Curran, P. M. Sharples, C. White, M. Knapp, "Time costs of caring for children with severe disabilities compared with caring for children without disabilities", *Developmental Medicine and Child Neurology*, 43, 529-533 (2001).
- [4] Y. Demiris, "Prediction of intent in robotics and multi-agent systems", *Cognitive Processing*, 8: 151-158, 2007.
- [5] J. Nicholson and M. Bonsall, "Powered Mobility for Children under Five Years of Age in England", *British Journal of Occupational Therapy*, 65(5):291-293 (2002).
- [6] J. Furumasu, P. Guerette and D. Tefft, "Relevance of the Paediatric Powered Wheelchair Screening Test for children with cerebral palsy", *Developmental Medicine and Child Neurology*, 46: 468-474 (2004)
- [7] M. A. Jones, I. R. McEwen and L. Hansen, "Use of Powered Mobility for a Young Child with Spinal Muscular Atrophy", *Physical Therapy*, 83(3): 253-262 (2003).
- [8] L. Fehr, W. E. Langbein, and S. B. Skaar, "Adequacy of power wheelchair control interfaces for persons with several disabilities: A clinical survey", *Journal of Rehabilitation Research and Development*, 37(3): 353-360 (2000).
- [9] M. Bottos, C. Bolcati, L. Sciuto, C. Ruggeri, and A. Feliciangeli, "Powered wheelchairs and independence in young people with tetraplegia", *Developmental Medicine and Child Neurology*, 43:769-777 (2001).
- [10] J. Dawson and H. Thornton, "Can Patients with Unilateral Neglect following Stroke drive Electrically Powered Wheelchairs?", *British Journal of Occupational Therapy*, 66(11):496-504 (2003).
- [11] A. Davies, L. H. De Souza, and A. O. Frank, "Changes in the quality of life in severely disabled people following provision of powered indoor/outdoor chairs", *Disability and Rehabilitation*, 25(6):286-290 (2003).
- [12] T. Carlson and Y. Demiris, "Using Visual Attention to Evaluate Collaborative Control Architectures for Human Robot Interaction", Proceedings of the AISB HRI Symposium, Edinburgh, Scotland, to appear (2009).