

Towards Using Social HRI for Improving Children’s Healthcare Experiences

Mary Ellen Foster

School of Computing Science
University of Glasgow
Glasgow, Scotland, United Kingdom
MaryEllen.Foster@glasgow.ac.uk

Ronald P. A. Petrick

Edinburgh Centre for Robotics
Heriot-Watt University
Edinburgh, Scotland, United Kingdom
R.Petrick@hw.ac.uk

Abstract

This paper describes a new research project that aims to develop a social robot designed to help children cope with painful and distressing medical procedures in a clinical setting. While robots have previously been trialled for this task, with promising initial results, the systems have tended to be teleoperated, limiting their flexibility and robustness. This project will use epistemic planning techniques as a core component for action selection in the robot system, in order to generate plans that include physical, sensory, and social actions for interacting with humans. The robot will operate in a task environment where appropriate and safe interaction with children, parents/caregivers, and healthcare professionals is required. In addition to addressing the core technical challenge of building an autonomous social robot, the project will incorporate co-design techniques involving all participant groups, and the final robot system will be evaluated in a two-site clinical trial.

Introduction

Children regularly experience pain and distress in medical situations that can produce both short-term (e.g., fear, distress, inability to perform procedures) and long-term (e.g., needle phobia, anxiety) negative effects (Stevens et al. 2011). While a range of techniques have been demonstrated as effective methods for managing such situations (e.g., breathing exercises, distraction techniques, cognitive-behavioural interactions (Chambers et al. 2009)), delivered through a variety of means (e.g., distraction cards, kaleidoscopes, music, and virtual reality games), recent studies have also explored the use of social robots as a tool for managing child pain and distress during medical procedures (Ali et al. 2019; Trost et al. 2019).

One significant drawback in many of these past studies is that the robots tend to be teleoperated or use purely scripted behaviours with limited autonomy and responsiveness, thereby reducing the overall flexibility and robustness of the system. Such approaches also limit the ability to personalise or adapt the system to different user groups (e.g., children, parents of children undergoing medical pro-



Figure 1: A child interacting with the Nao robot.

cedures, and medical practitioners), in order to provide appropriate context-dependent behaviours and responses.

From a technical point of view, a fundamental component in any social robot is the action selection system which controls overall robot behaviour: the robot must make high-level decisions as to which spoken, non-verbal, and task-based actions should be taken next by the system as a whole. It is also crucial not only to choose the appropriate action, but also to monitor the execution of such actions and the state of the world detected by the robot’s sensors: particularly in the context of embodied interactions with a robot, it is likely that the predicted state will often differ from the sensed world state, due to both the unexpected behaviour of the human interaction partners as well as the inherent uncertainty involved in sensing and acting in the physical world.

This paper describes a new project (Foster et al. 2020) aimed at addressing this limitation by developing and evaluating a clinically relevant and responsive social robot, using the Nao robot platform from SoftBank Robotics (Figure 1). While the majority of social robotics systems generally use either scripted behaviour for action selection, or machine learning approaches to learn the correct responses to user actions given sample inputs, we will instead use epistemic planning techniques (Dissing and Bolander 2020; Petrick and Foster 2013) as a key aspect of our approach, as the basis for high-level action selection and execution monitoring (Petrick and Foster 2020). In the remainder of this paper we outline the relevant related work, proposed technical architecture, and challenges on this project.

Related Work

The area of socially assistive robotics (Feil-Seifer and Mataric 2005) generally focuses on designing a robot system with the goal of creating an effective interaction with a human partner for the purpose of providing assistance and achieving measurable progress in a defined domain. Such robots have been used to improve the cognitive abilities of adult Alzheimer’s patients (Tapus, Tapus, and Mataric 2009), to alleviate feelings of loneliness and depression in the elderly (Wada et al. 2004), and to help adults with autism to improve work-related social skills (McKenna et al. 2019). An important application of socially assistive robots also focuses on autism in children, where robots have been used for diagnosis, intervention, and therapy (Cabibihan et al. 2013). Since children often perceive social robots as being similar to a companion animal or pet, such robots have also been used for play therapy and social learning (Breazeal 2011).

Our project will explore the use of socially assistive robots for reducing child distress and pain in clinical settings. (Trost et al. 2019) recently examined studies where a robot was used in this context: overall, while the results seem promising and suggest that the robots succeeded in reducing pain, a need for improved methodology and measures was identified. In particular, the authors suggest more effective approaches could be created by ensuring healthcare experts and system engineers collaborate from the start, and that user and family partners contribute to a user-centred design process. Our planned work includes input from all such groups as part of the research team (see below).

On the technical side, planning and interaction have a long history, and planning techniques have been applied previously in a range of social robots and interactive systems—recent examples include (Waldhart, Gharbi, and Alami 2016; Sanelli et al. 2017; Kominis and Geffner 2017; Papaioannou, Dondrup, and Lemon 2018). The most similar approach to ours is the JAMES social robot bartender (Petrick and Foster 2013; 2020), which directly used an automated planner to choose the robot’s physical, sensing, and interactive actions. This system will form the basis of the approach used on this project. Recent work on explainable planning (Fox, Long, and Magazzeni 2017) has also highlighted the links between planning and user interaction, and is relevant to this work.

Robot Behaviour and Task Environment

The high-level goal of the project is to develop and evaluate an autonomous social robot designed to help children deal with procedural pain in emergency room environments. The behaviour of the robot will be based on existing cognitive behavioural interventions that have been demonstrated to be effective in this context (e.g., distraction, empathy). The target robot platform is the Nao robot from SoftBank Robotics (Figure 1), which has been widely used in child-robot interaction studies, including in the identical clinical context we are targeting (Ali et al. 2019). In addition to lab-based testing with the robot, the system will be tested in the target environment throughout the project period, culminating in a two-site randomised clinical trial at the end of the project.

At the action level, the robot will engage in physical, sensory, and social actions. Physical actions will include movement of the robot as a whole (movement between locations, using build-in features like dancing) and moving various aspects of the embodiment (e.g., arm and head movements). Sensory actions will include targeted information gathering through available sensor modalities (vision, speech) about the social context and task, but also through verbal interaction (question asking). Social actions will include a range of dialogue and interactive actions (question answering, engagement behaviour, explanation). Moreover, the execution of some actions may change depending on the human user involved and the social context (e.g., answering a question for a parent versus a medical professional, interacting with a happy child versus one who is crying). The following robot utterances have been used in previous studies (Ali et al. 2019; Trost et al. 2020) and are also relevant here:

- “I am excited to play with you today.”
- “I need to catch my breath and control my breathing. Let’s try this together! Breathe in through your nose, hold your breath for two seconds, and out through your mouth. Try it with me!”
- “I understand you are in a lot of pain right now and I’m here to help you.”
- “I think we should celebrate how brave you are. Here’s a new dance I just learned.”

While the robot hardware platform itself supports a range of built-in features, which can be abstracted into high-level behaviours to support the needed actions, the actual decision as to which behaviours are included will be decided as part of a co-design process that involves the participation of children, parents, healthcare professionals, as well as the technical research team. This process therefore has implications for the technical aspects of the system design, including the process of modelling such actions for use with the action selection and decision-making components of the system.

Overview of the Robot System

The robot system will include components for social signal processing, high-level behaviour selection, and execution monitoring and recovery as shown in Figure 2. The system will be based on the JAMES robot bartender system (Petrick and Foster 2013), which uses a knowledge-level planner to generate plans for social interaction.

Social signal recognition: A core task in the system is to use the information from the robot’s built-in audio-visual sensors (possibly combined with environmental sensors) to determine the state of people (children, adults) in the task area. The particular states to be detected will be determined through a combination of the capabilities of the sensors, as well as the states determined to be relevant from the design process. Based on the detected verbal and non-verbal cues employed by all humans near the robot, the state will be estimated, including the physical state and social aspects such as attitudes, emotions, and intention. We plan to use a neural-network approach to detect the states, similar to the approach of (Roffo et al. 2019).

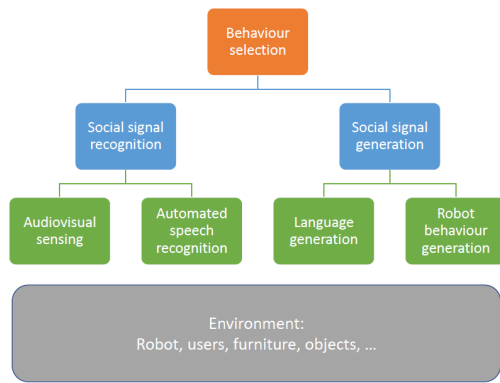


Figure 2: Software architecture for the robot system.

Behaviour selection: Using the available robot actions and the detected social states from the social signal recogniser, the system will choose appropriate high-level actions to be performed by the robot. Action selection is performed using PKS (Planning with Knowledge and Sensing) (Petrick and Bacchus 2002; Petrick and Foster 2013), an epistemic planner that operates at the knowledge level and reasons about how its knowledge state, rather than the world state, changes due to action. PKS has been previously used in interactive environments, such as the JAMES robot bartender, where information-gathering dialogue actions were modelled as sensing actions in the planner (Petrick and Foster 2020).

Social signal generation: Once a behaviour is selected, the social signal generation process generation converts these high-level actions into concrete robot-level actions that can be executed on the Nao robot platform. As output, this component will produce multimodal behaviour plans, including verbal and non-verbal actions, that are coordinated temporally and spatially for robot execution.

Execution monitoring and recovery: The system will monitor the changes to the state (as detected by the social signal recogniser) while planned actions are executed on the robot, using traditional plan monitoring techniques, but involving the richer social and epistemic states afforded by the state recognition and planning processes. Due to the inherent uncertainty of robot sensors and the unpredictable behaviour of humans it is likely that the planned states will often differ from the actual world state. The monitoring system will detect such mismatches and determine whether the execution of the current high-level plan should continue or whether a new plan is needed, invoking replanning as needed.

All software components will be developed using the Robot Operating System (ROS) as the core middleware, together with ROSPlan (Cashmore et al. 2015) framework for integrating the planning tools.

Challenges

From a high-level planning perspective, the task of applying planning in a child-centred medical context centres around an important knowledge engineering task to accurately mod-

elling the required states, actions, and goals that reflect the types of activities the robot is expected to perform. The high-level planner is responsible for selecting robot actions to respond appropriately in the current social state of the task, with a mix of physical, sensory, and social behaviours. However, unlike many robot systems which are designed from the ground up by the technical/research team, this project involves a wider collaboration at the design phase. Furthermore, the output of the robot may have to be tailored to different user groups, depending on the interaction situation. Finally, the entire robot system needs a high degree of robustness since the user studies will include a full clinical trial. We highlight some of these challenges below.

Co-design of robot capabilities: At a high level, the robot must select appropriate actions to support children undergoing clinical procedures. However, the details of the exact behaviours and features will be developed through a co-design approach using the principles of user-centred interaction design (Preece, Sharp, and Rogers 2015). The perspectives of all groups involved in the task will be considered, including children, parents/caregivers, healthcare professionals, as well as the research and technical team. This process will consider all aspects of the task including the needs of children, caregivers, and the medical team during a clinical procedure, the (possibly differing) perceptions these groups have of the robot, and the core functionality required for the task. While this process will be somewhat constrained by the physical limitations of the robot platform, and the representational and reasoning capabilities of the planner, the knowledge engineering task of building planning models cannot be fully realised without the input from the co-design phase.

Personalised interaction: The robot system is meant to interact with different user groups, from children to healthcare professionals. The kinds of plans that are generated for the various groups will also differ, as will the potential execution of those plans on the robot. For instance, answering a child's question may be realised differently compared with a similar question from a parent or healthcare professional. Some of these differences will be reflected at the action level in the planning domain model, which also introduces some interesting connections to explainable planning (XAIP) (Fox, Long, and Magazzeni 2017). In particular, explainability of planning decisions and plan execution will be explored on the project, especially in the context of analysing planner decisions by healthcare professionals.

System robustness and safety: Several user studies are planned for the project, from initial lab-based testing to a clinical trial, to understand the behaviour of the system. The end system is meant to be run in a real clinical setting (at least for evaluation purposes), with minimal human intervention. Moreover, the system is meant to run in the presence of different user groups, including both children and adults. As a result, the system will require a high degree of safety and robustness at all levels. In particular, recognising children's social signals in a real-world setting is expected to present a particular technical challenge, and we will explore a range of sensor configurations and audiovisual processing approaches to find the one that works best in this

specific domain. In the final phase of the project, the clinical trial will evaluate the primary goal of the project: that interaction with a robust, adaptive, socially intelligent robot can effectively support children during a clinical procedure and reduce their distress and pain. The clinical trial is planned for the two Canadian paediatric departments where the co-design and usability studies will also take place.

Summary and Conclusions

This paper describes a new research project that plans to design and use a social robot in a medical setting to support children undergoing painful clinical procedures. At the core of the action selection and decision-making system will be a high-level epistemic planner that will be responsible for generating plans to interact with different user groups, including children, parents/caregivers, and healthcare professionals. The robot behaviour will include physical robot actions, sensing actions, and social actions for guiding the interaction. A co-design process will be used, involving all the user groups along with the research/technical team, which will feed into the design of the planning model. Some aspects of explainable planning and plan personalisation will also be necessary on the project. The system will be demonstrated and evaluated through a series of usability studies, leading up to a full clinical trial in a real-world medical setting. The hope is that this work will extend existing work on social robotics and demonstrate the potential of automated planning as a useful tool for autonomous decision making in this challenging domain.

Acknowledgements

This work is funded by the ESRC/SSHRC Canada-UK Artificial Intelligence Initiative through grant ES/T012986/1. We also acknowledge the lead Canadian research partners on the project: Samina Ali, Sasha Litwin, Jennifer Parker, David Harris Smith, Jennifer Stinson, and Frauke Zeller.

References

- Ali, S.; Manaloor, R.; Ma, K.; Sivakumar, M.; Vandermeer, B.; Beran, T.; Scott, S.; Graham, T.; Curtis, S.; Jou, H.; Beirnes, N.; and Hartling, L. 2019. LO63: Humanoid robot-based distraction to reduce pain and distress during venipuncture in the pediatric emergency department: A randomized controlled trial. *Canadian Journal of Emergency Medicine* 21(S1):S30–S31.
- Breazeal, C. 2011. Social robots for health applications. In *Int. Conf. of the IEEE Engineering in Medicine and Biology Society*.
- Cabibihan, J.-J.; Javed, H.; Ang, Marcelo, J.; and Aljunied, S. M. 2013. Why robots? A survey on the roles and benefits of social robots in the therapy of children with autism. *International Journal of Social Robotics* 5(4):593–618.
- Cashmore, M.; Fox, M.; Long, D.; Magazzeni, D.; Ridder, B.; Carrera, A.; Palomeras, N.; Hurtos, N.; and Carreras, M. 2015. ROS-Plan: Planning in the Robot Operating System. In *ICAPS*.
- Chambers, C. T.; Taddio, A.; Uman, L. S.; and McMurtry, C. 2009. Psychological interventions for reducing pain and distress during routine childhood immunizations: A systematic review. *Clinical Therapeutics* 31:S77–S103.
- Dissing, L., and Bolander, T. 2020. Implementing Theory of Mind on a Robot Using Dynamic Epistemic Logic. In *IJCAI*, 1615–1621.
- Feil-Seifer, D., and Mataric, M. 2005. Socially assistive robotics. In *IEEE Int. Conf. on Rehabilitation Robotics*.
- Foster, M. E.; Ali, S.; Litwin, S.; Parker, J.; Petrick, R. P. A.; Smith, D. H.; Stinson, J.; and Zeller, F. 2020. Using AI-enhanced social robots to improve children’s healthcare experiences. In *The 12th International Conference on Social Robotics (ICSR 2020)*.
- Fox, M.; Long, D.; and Magazzeni, D. 2017. Explainable planning. In *Proceedings of the IJCAI Workshop on Explainable AI*.
- Kominis, F., and Geffner, H. 2017. Multiagent online planning with nested beliefs and dialogue. In *International Conference on Automated Planning and Scheduling*.
- McKenna, P.; Broz, F.; Keller, I.; Part, J. L.; Rajendran, G.; and Aylett, R. 2019. Towards robot-assisted social skills training for adults with ASC. In *CHI 2019 Workshop on the Challenges of Working on Social Robots that Collaborate with People*.
- Papaoiannou, I.; Dondrup, C.; and Lemon, O. 2018. Human-robot interaction requires more than slot filling - multi-threaded dialogue for collaborative tasks and social conversation. In *FAIM/ISCA Workshop on AI for Multimodal Human Robot Interaction*, 61–64.
- Petrick, R. P. A., and Bacchus, F. 2002. A knowledge-based approach to planning with incomplete information and sensing. In *Proceedings of AIPS 2002*, 212–221.
- Petrick, R. P. A., and Foster, M. E. 2013. Planning for social interaction in a robot bartender domain. In *International Conference on Automated Planning and Scheduling*, 389–397.
- Petrick, R. P. A., and Foster, M. E. 2020. Knowledge engineering and planning for social human-robot interaction: A case study. In Vallati, M., and Kitchin, D., eds., *Knowledge Engineering Tools and Techniques for AI Planning*. Springer. 261–277.
- Preece, J.; Sharp, H.; and Rogers, Y. 2015. *Interaction design: beyond human-computer interaction*. John Wiley & Sons.
- Roffo, G.; Vo, D.-B.; Tayarani, M.; Rooksby, M.; Sorrentino, A.; Di Folco, S.; Minnis, H.; Brewster, S.; and Vinciarelli, A. 2019. Automating the administration and analysis of psychiatric tests. In *CHI Conference on Human Factors in Computing Systems*.
- Sanelli, V.; Cashmore, M.; Magazzeni, D.; and Iocchi, L. 2017. Short-term human-robot interaction through conditional planning and execution. In *International Conference on Automated Planning and Scheduling*.
- Stevens, B. J.; Abbott, L. K.; Yamada, J.; Harrison, D.; Stinson, J.; Taddio, A.; Barwick, M.; Latimer, M.; Scott, S. D.; Rashotte, J.; Campbell, F.; and Finley, G. A. 2011. Epidemiology and management of painful procedures in children in Canadian hospitals. *Canadian Medical Association Journal* 183(7):E403–E410.
- Tapus, A.; Tapus, C.; and Mataric, M. J. 2009. The use of socially assistive robots in the design of intelligent cognitive therapies for people with dementia. In *IEEE Int. Conf. on Rehabilitation Robotics*.
- Trost, M. J.; Ford, A. R.; Kysh, L.; Gold, J. I.; and Matarić, M. 2019. Socially assistive robots for helping pediatric distress and pain. *The Clinical Journal of Pain* 35(5):451–458.
- Trost, M. J.; Chrysilla, G.; Gold, J. I.; and Matarić, M. 2020. Socially-assistive robots using empathy to reduce pain and distress during peripheral IV placement in children. *Pain Research & Management* 2020:1–7.
- Wada, K.; Shibata, T.; Saito, T.; and Tanie, K. 2004. Effects of robot-assisted activity for elderly people and nurses at a day service center. *Proceedings of the IEEE* 92(11):1780–1788.
- Waldhart, J.; Gharbi, M.; and Alami, R. 2016. A novel software combining task and motion planning for human-robot interaction. In *AAAI Fall Symposium Series*.