Using AI Planning for Managing Affective States in Social Robotics

Alan Lindsay
Heriot-Watt University
Edinburgh, UK

Andrés A. Ramírez-Duque
University of Glasgow
Glasgow, UK

Mary Ellen Foster
University of Glasgow
Glasgow, UK

Ronald P. A. Petrick
Heriot-Watt University
Edinburgh, UK

ABSTRACT

Social robotics has recently focused on developing AI agents that recognise and respond to human emotions. The use of plan-based approaches is promising, especially in domains where collecting data in advance is challenging (e.g., medical domains). However, we observe that the appropriate use of the user' affective state will vary with the particular interaction, the expected impact of the robot's behaviours on the user, and the opportunity and accuracy of affective sensing. We observe that there are different ways of modelling the user's affective state, and the appropriate choice will take into consideration the relationship between the user's affective state and the robot's behaviour. We propose alternative methods of modelling the user's affective state, and use lessons learnt from a recent project in order to discuss the relevant factors in each approach. We use simulated data in order to demonstrate the flexibility of model-based generation of interaction strategies.

CCS CONCEPTS

• **Human-centered computing → HCI theory, concepts and models**; • **Computer systems organization → Robotic autonomy**; • **Computing methodologies → Planning under uncertainty**.

KEYWORDS

Managing affective state, Plan-based interaction, Socio-Affective sensing, Human-Robot interaction

ACM Reference Format:
Alan Lindsay, Andrés A. Ramírez-Duque, Mary Ellen Foster, and Ronald P. A. Petrick. 2024. Using AI Planning for Managing Affective States in Social Robotics. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3610978.3640744

1 INTRODUCTION

Recent advances in the fields of Human-Computer and Human-Robot Interaction (HRI) have focused on developing AI agents that recognise and respond to human emotions [28]. These approaches take advantage of the notion that emotions provide a high-level evaluation of events based on their relationship with human internal beliefs and desires. In particular, emotions provide another communication modality to interpret an individual's implicit cues and needs [23]. As a consequence, using socio-affective cognitive models has improved human-robot interaction in collaborative tasks requiring socially acceptable strategies and in some cases can (indirectly) improve harmful behaviours related to emotions such as fear, stress, or anxiety in mental health care [5].

Typically, researchers assume that the robot is equipped with a social signal processing system capable of estimating basic and explicit affective expressions. These signals have been used to reactively generate sequences of robot behaviours, with the aim of communicating the robot's mental state and communicative intentions [1, 21, 22]. Within deliberative architectures, monitoring of user's affective state has been used to adapt actions in human-robot collaborative tasks [10] and to provide a social robot with skills to explain its internal reasoning process [21]. However, deliberative reasoning that directly reasons with specific aspects of the user's affective state to manage complex emotions such as anxiety has not been fully explored. Plan-based approaches are promising as they support the creation of environment models, which can encode interaction and social constraints while also embedding appropriate domain knowledge.

In this work, we consider the management of the user's affective state. In particular, we consider how to incorporate the user's affective state into a planning model that underpins a human-robot interaction. As a running example, we use our recent experience of building a companion robot for a medical scenario. Using this example as grounding, we discuss the use of deterministic and non-deterministic approaches to modelling the user's affective state. We demonstrate how a planning model can be used to capture various strategies that can be specialised both for a particular user (e.g., in reaction to higher-than-expected anxiety) or a particular supporting strategy (e.g., focusing on either diverting behaviours or adopting cognitive behavioural interventions).

2 BACKGROUND

Socially assistive robots (SARs) are embodied devices designed to interact with humans by communicating through mechanisms compatible with a human-centric approach [7]. The primary focus of SARs is to provide necessary aid to humans by engaging with them socially. Research has shown that SARs can help alleviate tension,

This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0323-2/24/03.
https://doi.org/10.1145/3610978.3640744

reduce stress, and enhance social interactions in medical settings. In addition, SARs have proven helpful by assisting and supporting people experiencing stress or anxiety, such as children undergoing medical procedures [6, 15, 24]. Studies have compared short-term single-procedure exposure or long-term companionship [11], as well as the effectiveness of robot-delivered interventions, such as distraction and cognitive-behavioural therapies, with standard care in needle-based practices [2, 19, 20].

2.1 Running Example: A Companion Robot for a Medical Scenario

This work uses a running example involving supporting children during a painful and distressing medical procedure. A robot is placed in a small room together with the patient, their carers and a Health Care Provider (HCP) during the course of a single clinical procedure. In related work [8, 12, 13], we have developed a fully functioning companion robot for operating in this scenario, which was designed using both a co-design strategy (involving several stages and children, parents and HCPs) and targeted meetings between the technical team and the HCPs. IV insertion was identified as an appropriate procedure: This is one of the most commonly performed procedures in the context of children seeking medical care. A standard procedure in a paediatric setting is to provide distraction before and during the procedure to alleviate pain and distress. We identified several stages: introduction, preprocedure (optional site-check), procedure, debrief, and conclusion. The robot positions itself as a friendly and supportive companion, setting out positive expectations, and can present various supportive behaviours, including humour, coping strategies, role modelling, and positive reinforcements.

2.2 Planning Based Interaction

Planning as a high-level decision-making mechanism has previously been integrated with cognitive models in various social robots and interactive systems, e.g., [14, 18], but without estimating and managing the user's affective state. Managing of the user's affective state has been considered in planning models [26, 27], but these approaches target specific aspects of social and affective state and do not present generally applicable approaches.

In our approach, the robot's behaviours are underpinned by a planning model, which uses a declarative representation to encode the domain knowledge and possible interactions concisely. We use a fully observable non-deterministic (FOND) planning model based on [16], which can be defined as a tuple $\langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$, with fluents $\mathcal{F}$, initial state $\mathcal{I}$ (a full assignment to $\mathcal{F}$), a partial goal state $\mathcal{G}$, and a set of actions $\mathcal{A}$. Each action $a \in \mathcal{A}$ is a pair $\langle pre_a, eff_a \rangle$, with a precondition pre_a (a subset of $\mathcal{F}$ that must hold) and an effect eff_a (a set of possible outcomes—fluents that are made true or false). If an action defines one outcome, it is a deterministic action; otherwise, it is a non-deterministic action. Each action application results in an outcome, but the outcome cannot be chosen by the planner. A solution to the problem is a branched plan π, which includes alternative action outcomes and describes the sequence of actions that will achieve the goal, given any outcome. When we can assume that all actions have a single outcome (a deterministic model), a plan can be represented as a sequence of actions. In deterministic models we also allow numeric fluents in the state, which can be used in the preconditions and effects of actions, and can be used in the definition of the optimisation function.

3 MODEL-BASED MANAGING OF THE USER'S AFFECTIVE STATE

A key problem in designing SAR interactions is in managing the user's perception of those events. This relies on careful design of the interaction, making changes to the strategy where the user reacts differently than expected and exploiting the shared context of the unfolding events. As this varies with the specific relationships between the users and the environment, there is no general approach to incorporating all aspects of the user's affective state into an existing model. However, through observations and through consultation with experts, we can build the monitoring and management of affective state into the interaction.

Based on a recent co-design study [8], specifically targeted to our running example, requirements were identified that covered a broad range of aspects, including the robot's appearance, the interaction, and the ethical considerations. Using this information we have developed a fully working system, including social signal recognition, planning, and robot behaviour rendering modules, implemented on embedded hardware. The system uses a planning model to underpin the robot's interaction. The main interaction captured in the planning model is structured along the possible patient pathways outlined with HCPs during the design process. The robot can perform a range of actions, each of which is represented in the planning model: distracting actions (e.g., dancing) and calming and instructive actions (e.g., breathing exercises); sensing actions for the medical scenario, e.g., to maintain the progress through the medical procedure; and patient focused sensing actions, e.g., to determine whether the patient is engaged in the interaction. The main stages of the medical procedure (see Subsection 2.1) were used to organise the appropriate behaviours and in order to specify key objectives for the robot in each stage. For example, we can ensure that the robot delivers certain key information to the patient during the preprocedure (e.g., regarding its role).

In this section we examine two alternative approaches for modelling the user's affective state within the system's planning model. We also provide an overview of our sensing approach and discuss how the limitations in the sensing component influence the selection of an appropriate planning model.

3.1 Modelling the User's Affective State Changes in Action Effects

A basic requirement for integrating affective state into an interaction is that the interaction is conditioned on the value of the affective state. We can assume that there is a finite set of aspects of the user's affective state that are relevant to the interaction, and we associate each of these aspects with a variable, v. Each variable v has a set of possible values, $\mathbb{D}(v)$, which captures the set of discrete values for v. For example, in the context of anxiety, this might be a pair, e.g., $\mathbb{D}(v_{anxiety}) = (\text{'}OK\text{'}, \text{'}high\text{-}anxiety\text{'})$, or a set of values: $\mathbb{D}(v_{anxiety}) = (1(low), 2, 3, 4, 5(high))$, depending on the level of accuracy required in the model.

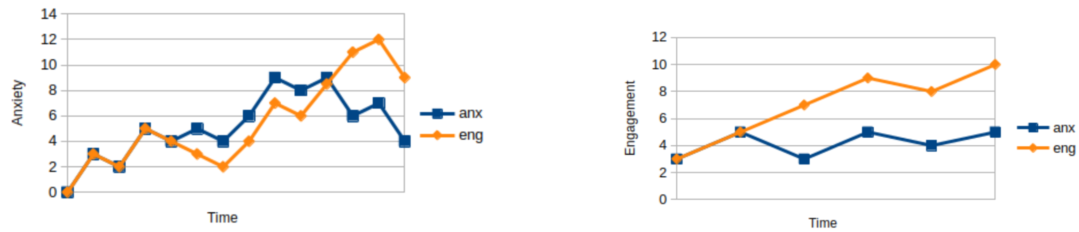

(a) Anxiety for simulated plans

(b) Engagement for simulated plans[1]

Figure 1: Two plans were generated, the first to reduce anxiety (*anx* in blue), and the second to increase engagement (*eng* in orange). These plots present the plan values of anxiety (a) and engagement (b) against time.

Action	Anxiety	Action	Anxiety
`ready_to_start`	+3	`start_procedure`	
`introduction`	-1	*with procedure plan*	+2
`divert dance1`	-2	*no plan*	+5
`educate ivdescription`	+1	`imp_plan breathing`	
`make_procedure_plan`	0	*practised breathing*	-2
`debrief song1`	-3	*not practised*	0

Table 1: Predicted impact on user anxiety for a selection of the planning actions interaction model for a medical scenario.

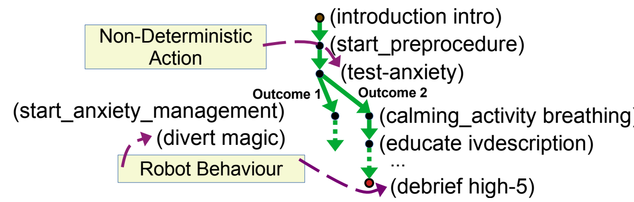

Figure 2: A partial plan showing actions (e.g., breathing exercises and high five), sensing actions (e.g., testing patient anxiety) and procedure actions (e.g., the preprocedure start).

To incorporate a representation of affective state change into a planning model, we assess each of the transitions in the model and assess their expected impact on the user's affective state. The main advantage of representing the impact of transitions on the affective state within the model is that the planner can reason about its value directly. For example, Table 1 presents values indicating how user anxiety is predicted to be impacted by aspects of the medical procedure (e.g., the `start_procedure` action), and the robot's intervention (e.g., `divert dance1`). Notice that past aspects of the interaction can impact on the predicted effect of behaviours on anxiety level (e.g., the anxiety for the procedure depends on whether or not the robot has made a plan of how to support the patient during the procedure). Given these values for expected changes on anxiety and an objective function (e.g., minimise anxiety, or maintain an acceptable level of anxiety), the planner can search for a sequence of actions that optimises/satisfies this objective.

We compiled these anxiety values into the planning model and used it to generate two plans (using the Optic planner [3]), each optimising a different metric. In [9] they demonstrated two distinct strategies that the robot could use. The first was to use cognitive-behavioural interventions, e.g., the robot might practice breathing strategies with the patient during the preprocedure, to help the patient stay calm during the procedure stage. An alternative strategy was to focus on distracting the patient (e.g., dancing or a quiz). We used this as motivation for two optimisation functions: The first aims to minimise anxiety, whereas the second aims to maximise engagement. In Figure 1 we have plotted each of the resulting plans against both metrics. We can see from the engagement plot (1b) that the engagement plan (orange line) is focused on increasing engagement, and initially is effective at managing anxiety. However, the lack of preparation during the preprocedure leads to higher anxiety during the procedure stage (see 1a). This strategy involved selecting diverting actions at each opportunity. Conversely, in the case of anxiety (blue line) the anxiety plot (1a) demonstrates that the preparation during the preprocedure (practising breathing, making a plan, and providing some information to the patient), have led to lower simulated anxiety during the procedure. As a consequence, the engagement plot indicates that the interaction was not as engaging. It is interesting to note that these two plans, with quite different strategies (corresponding to the distraction and cognitive-behavioural strategies used in [9]) were generated from a single planning model (with different metric functions).

In order for the model to be useful, and progress the interaction towards its goals, it is important that the model should be accurate. For example, in [25] they indicate the importance of including the therapist in their modelling loop (they were creating personalised physical rehabilitation plans). In future work we will investigate how continual monitoring can be used so that variations between the predicted state and the sensed environmental state can be detected and replanning used [4].

3.2 Incorporating Affective Sensing

In general, accurately modelling affective state within the model can be impractical due to limited data, the potential complexity of the relationship, and the lack of relevance of the planning model to the important aspects affecting affective state. In our medical scenario, a key concern for the HCP during the procedure is that patients will typically harbour anxiety inducing negative anticipation and uncertainty about the procedure. The system, therefore, makes a

[1]The engagement values have been modelled to support optimisation.

plan of how the robot will support the child during the procedure and provides the child with certainty. The aim is to reduce their negative anticipation and uncertainty by replacing it with positive anticipation. We could attempt to capture this intervention's impact on anxiety through observations. However, it is likely that the impact of this strategy will be fairly subtle across the interaction, and without expert input, it would be challenging to model as an effect without substantial data.

As an alternative, we can consider incorporating sensing into the model, allowing certain outcomes' uncertainty to be explicitly represented. We define a sensing action for each variable, v, with an outcome for each of the $|\mathbb{D}(v)|$ possible values, such that outcome i adds a proposition, p_i^v, representing the i^{th} value of v. The action preconditions use these propositions to identify the required affective state variables, and their values.

Figure 2 presents part of a branched plan (generated using the PRP planner [16]), including the sensing action `sense_anxiety`. This figure demonstrates how branched plans captures the strategy in the context of a variety of patient pathways, customising the interaction based on the specific detail of the medical procedure, and the patient's affective state and preferences. This allows the plan to capture the strategy in the case of either high (e.g., selecting an appropriate intervention) or normal anxiety (e.g., continuing the interaction by practising breathing exercises). In our medical scenario, we define a variable for engagement with values: $\mathbb{D}(v_{engaged}) = (\text{'}OK'\text{'}, \text{'}low\text{-}engagement'\text{'})$. This value is monitored at certain points during the procedure, allowing the robot to adapt its strategy. For example, in the context of high patient anxiety, the value $v_{engaged}$ is monitored. In the case study, in cases where the patient is not engaging, the robot stops its interaction.

Anxiety Management Component. As part of our design process, we have identified specific points during the interaction where there are appropriate interventions for managing the patient's affective state. To exploit these opportunities, we pair targeted sensing of patient's affective state with managing interventions. In this part we focus on a specific anxiety management intervention during the preprocedure. The anxiety test action (a sensing action that determines whether the patient's anxiety level is OK) is used as a sensing action, and in cases where the patient has high anxiety, we adopt the anxiety management intervention.

Our approach is based on first detecting high anxiety and subsequently enacting a specially constructed intervention sequence, which was designed with HCPs. The final step involves retesting, which is key to allowing the interaction to proceed appropriately.

An example intervention adopts a strategy of first diverting with a high-distraction activity, such as dancing, and then attempting to calm the patient using a relaxing activity, such as breathing exercises. The belief is that the patient's anxiety will be reduced by this procedure. This belief is confirmed by retesting the patient's anxiety level at the end of the procedure. If the patient's anxiety has not reduced, an appropriate subsequent course of action should be chosen. In the case of our scenario, in order to minimise further distress, the robot will step back from the interaction.

Monitoring Affective State in a Medical Setting. The development of a planning model for the management of affective state must be made in parallel with an investigation of the sensing opportunities in the scenario and the accuracy that can be achieved. In the context of affective state prediction, these models typically involve the design of multi-stage pipelines, which incorporate information from multiple sensory sources in order to reduce ambiguities. However, due to the physical constraints of the robot deployment, the system was limited to facial analysis. In fact, predicting the patient's affective state automatically in this scenario proves to be challenging due to various limiting factors: 1.) There is limited space near the patient, and constant staff movement causes occlusion. 2.) The patient is likely to be wearing a surgical mask. 3.) The system must be portable and mobile; it must be able to move between the different emergency rooms with agility, which reduces the possibility of using fixed cameras and Internet connections via LAN and WLAN due to interference. With these limitations in mind, a pipeline has been developed to automatically analyse the patient's face in a medical scenario like the one described in our running example. The automatic facial analysis pipeline is based on Nvidia DeepStream SDK[17] and was deployed using a Jetson Nano board. During a practical application, six facial expressions, the focus of visual attention, and the speed of movement of the patient's head are estimated. The engagement ($v_{engaged}$) and the patient's affective state ($v_{anxiety}$) are then determined using the aforementioned social cues.

In our specific medical setting, we do not expect the sensing component to be able to distinguish more than two levels of anxiety. We are currently in the process of testing our system's usability using the non-deterministic model and two levels of anxiety. However, we hope to be able to support a richer model in other deployments where e.g., skin response can be used.

4 CONCLUSION AND FUTURE WORK

In this work, we considered the management of user affective state in plan-based socially assistive robots. We presented both deterministic and non-deterministic modelling approaches. In the deterministic approach, the impact on user affective state of robot behaviours and world events are captured within the planning actions, whereas in the non-deterministic model, the uncertainty of how the robot's interventions impact on affective state is captured explicitly within the planning model. We observed that while deterministic models require substantial data and expert interpretation, non-deterministic models allow for explicit testing of assumptions through sensing. We demonstrated that a benefit of deterministic models is that the associated planning systems are capable of effective optimisation, and we demonstrated how a single deterministic model can be used to generate alternative supporting strategies by using alternative metrics for optimisation. We then showed how alternative strategies for individual patient pathways are captured within the plans for the non-deterministic model.

ACKNOWLEDGMENTS

The authors wish to acknowledge the SSHRC-UKRI Canada-UK Artificial Intelligence Initiative (UKRI grant ES/T01296/1) for financial support of this project. We would also like to acknowledge all patient partners, research and clinical staff as well as the youth and families who have made this project possible.

REFERENCES

[1] Carole Adam, Wafa Johal, Damien Pellier, Humbert Fiorino, and Sylvie Pesty. 2016. Social Human-Robot Interaction: A New Cognitive and Affective Interaction-Oriented Architecture. In *Social Robotics*, Arvin Agah, John-John Cabibihan, Ayanna M. Howard, Miguel A. Salichs, and Hongsheng He (Eds.). Springer International Publishing, Cham, 253–263.

[2] S Ali, R Manaloor, K Ma, M Sivakumar, B Vandermeer, T Beran, S Scott, T Graham, S Curtis, H Jou, et al. 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* (2019).

[3] J Benton, Amanda Coles, and Andrew Coles. 2012. Temporal planning with preferences and time-dependent continuous costs. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 22. 2–10.

[4] Michael Cashmore, Andrew Coles, Bence Cserna, Erez Karpas, Daniele Magazzeni, and Wheeler Ruml. 2019. Replanning for situated robots. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 29. 665–673.

[5] Meia Chita-Tegmark and Matthias Scheutz. 2021. Assistive Robots for the Social Management of Health: A Framework for Robot Design and Human–Robot Interaction Research. *International Journal of Social Robotics* 13 (4 2021), 197–217. Issue 2. https://doi.org/10.1007/s12369-020-00634-z

[6] Julia Dawe, Craig Sutherland, Alex Barco, and Elizabeth Broadbent. 2019. Can social robots help children in healthcare contexts? A scoping review. *BMJ Paediatrics Open* 3, 1 (2019), e000371. https://doi.org/10.1136/bmjpo-2018-000371

[7] D. Feil-Seifer and M.J. Mataric. 2005. Socially Assistive Robotics. In *Proceedings of the International Conference on Rehabilitation Robotics*.

[8] Mary Ellen Foster, Patricia Candelaria, Lauren J. Dwyer, Summer Hudson, Alan Lindsay, Fareha Nishat, Mykelle Pacquing, Ronald P. A. Petrick, Andrés Alberto Ramírez-Duque, Jennifer Stinson, Frauke Zeller, and Samina Ali. 2023. Co-Design of a Social Robot for Distraction in the Paediatric Emergency Department. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 461–465. https://doi.org/10.1145/3568294.3580127

[9] Lindsay A. Jibb, Kathryn A. Birnie, Paul C. Nathan, Tanya N. Beran, Vanessa Hum, J. Charles Victor, and Jennifer N. Stinson. 2018. Using the MEDiPORT humanoid robot to reduce procedural pain and distress in children with cancer: A pilot randomized controlled trial. *Pediatric Blood & Cancer* 65, 9 (2018), e27242. https://doi.org/10.1002/pbc.27242

[10] Séverin Lemaignan, Mathieu Warnier, E. Akin Sisbot, Aurélie Clodic, and Rachid Alami. 2017. Artificial cognition for social human–robot interaction: An implementation. *Artificial Intelligence* 247 (2017), 45–69. https://doi.org/10.1016/j.artint.2016.07.002 Special Issue on AI and Robotics.

[11] Mike E.U. Ligthart, Mark A. Neerincx, and Koen V. Hindriks. 2022. Memory-Based Personalization for Fostering a Long-Term Child-Robot Relationship. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. 80–89. https://doi.org/10.5555/3523760.3523775

[12] Alan Lindsay. 2023. On Using Action Inheritance and Modularity in PDDL Domain Modelling. In *Proceedings of the International Conference on Automated Planning and Scheduling*.

[13] Alan Lindsay, Andrés Alberto Ramírez-Duque, Ronald P. A. Petrick, and Mary Ellen. Foster. 2022. A Socially Assistive Robot using Automated Planning in a Paediatric Clinical Setting. In *Proceedings of the AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction (AI-HRI)*.

[14] Francisco Martín, Francisco J. Rodríguez Rodríguez Lera, Jonatan Ginés, and Vicente Matellán. 2020. Evolution of a Cognitive Architecture for Social Robots: Integrating Behaviors and Symbolic Knowledge. *Applied Sciences* 10 (9 2020), 6067. Issue 17. https://doi.org/10.3390/app10176067

[15] Clara J Moerman, Loek van der Heide, and Marcel Heerink. 2018. Social robots to support children's well-being under medical treatment: A systematic state-of-the-art review. *Journal of Child Health Care* 23 (Nov 2018), 596–612. https://doi.org/10.1177/1367493518803031

[16] Christian Muise, Sheila McIlraith, and Christopher Beck. 2012. Improved non-deterministic planning by exploiting state relevance. In *Proceedings of the International Conference on Automated Planning and Scheduling*. 172–180.

[17] NVIDIA. 2022. NVIDIA DeepStream SDK. https://developer.nvidia.com/deepstream-sdk

[18] Ronald Petrick and Mary Ellen Foster. 2020. Knowledge engineering and planning for social human–robot interaction: A case study. In *Knowledge Engineering Tools and Techniques for AI Planning*.

[19] Silvia Rossi, Marwa Larafa, and Martina Ruocco. 2020. Emotional and Behavioural Distraction by a Social Robot for Children Anxiety Reduction During Vaccination. *International Journal of Social Robotics* 12, 3 (Jan 2020), 765–777. https://doi.org/10.1007/s12369-019-00616-w

[20] Matthijs H. J. Smakman, Koen Smit, Lotte Buser, Tom Monshouwer, Nigel van Putten, Thymen Trip, Coen Schoof, Daniel F. Preciado, Elly A. Konijn, Esther M. van der Roest, and Wouter M. Tiel Groenestege. 2021. Mitigating Children's Pain and Anxiety during Blood Draw Using Social Robots. *Electronics* 10, 10 (May 2021), 1211. https://doi.org/10.3390/electronics10101211

[21] Sonja Stange, Teena Hassan, Florian Schröder, Jacqueline Konkol, and Stefan Kopp. 2022. Self-Explaining Social Robots: An Explainable Behavior Generation Architecture for Human-Robot Interaction. *Frontiers in Artificial Intelligence* 5 (4 2022). https://doi.org/10.3389/frai.2022.866920

[22] Ana Tanevska, Francesco Rea, Giulio Sandini, Lola Cañamero, and Alessandra Sciutti. 2020. A Socially Adaptable Framework for Human-Robot Interaction. *Frontiers in Robotics and AI* 7 (10 2020). https://doi.org/10.3389/frobt.2020.00121

[23] Leimin Tian and Sharon Oviatt. 2021. A Taxonomy of Social Errors in Human-Robot Interaction. *ACM Transactions on Human-Robot Interaction* 10 (6 2021), 1–32. Issue 2. https://doi.org/10.1145/3439720

[24] Margaret J Trost, Adam R Ford, Lynn Kysh, Jeffrey I Gold, and Maja Matarić. 2019. Socially assistive robots for helping pediatric distress and pain: a review of current evidence and recommendations for future research and practice. *The Clinical journal of pain* (2019).

[25] Alessandro Umbrico, Marco Benadduci, Roberta Bevilacqua, Amedeo Cesta, Francesca Fracasso, Elvira Maranesi, Andrea Orlandini, and Gabriella Cortellessa. 2023. Combining Clinical and Spatial Constraints into Temporal Planning to Personalize Physical Rehabilitation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 33. 532–540.

[26] Loïs Vanhée, Laurent Jeanpierre, and Abdel-Illah Mouaddib. 2022. Anxiety-Sensitive Planning: From Formal Foundations to Algorithms and Applications. In *Proceedings of the International Conference on Automated Planning and Scheduling*.

[27] Zahra Zahedi, Mudit Verma, Sarath Sreedharan, and Subbarao Kambhampati. 2023. Trust-aware planning: Modeling trust evolution in iterated human-robot interaction. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 281–289.

[28] Raziyeh Zall and Mohammad Reza Kangavari. 2022. Comparative Analytical Survey on Cognitive Agents with Emotional Intelligence. *Cognitive Computation* 14 (7 2022), 1223–1246. Issue 4. https://doi.org/10.1007/s12559-022-10007-5