Kaspar Explains: The Effect of Causal Explanations on Visual Perspective Taking Skills in Children with Autism Spectrum Disorder*

Marina Sarda Gou[1], Gabriella Lakatos[1], Patrick Holthaus[1], Ben Robins[1], Sílvia Moros[1], Luke Wood[1], Hugo Araujo[2], Christine deGraft-Hanson[3], Mohammad Reza Mousavi[2], Farshid Amirabdollahian[1]

Abstract— **This paper presents an investigation into the effectiveness of introducing explicit causal explanations in a child-robot interaction setting to help children with autism improve their Visual Perspective Taking (VPT) skills. A sample of ten children participated in three sessions with a social robot on different days, during which they played several games consisting of VPT tasks. In some of the sessions, the robot provided constructive feedback to the children by giving causal explanations related to VPT; other sessions were control sessions without explanations. An analysis of the children's learning progress revealed that they improved their VPT abilities faster when the robot provided causal explanations. However, both groups ultimately reach a similar ratio of correct answers in later sessions. These findings suggest that providing causal explanations using a social robot can be effective to teach VPT to children with autism. This study paves the way for further exploring a robot's ability to provide causal explanations in other educational scenarios.**

I. Introduction

Autism Spektrum Disorder (ASD) is a neurodevelopmental disorder that affects mainly communication and social interaction skills. It is often characterized by the difficulty in establishing and maintaining relationships with peers, family members, and other individuals [1]. This can lead to isolation, frustration, and behavioural problems. Society does not provide sufficient provisions for such disorders and preparing children for engagement with society can have beneficial outcomes. Recent advancements in technology have opened up new opportunities for individuals with ASD to improve their social skills. In particular, humanoid social robots have shown promise as tools that can provide a controlled, safe and non-threatening environment where children with ASD can practice and enhance their social interaction and communication skills [2].

Visual Perspective Taking (VPT) skills are an important aspect of social interaction and communication. They relate to the ability to see the world from another person's perspective, taking into account what they see and how they see it [3]. VPT refers to a person's understanding that other people might have a different line of sight to themselves, and to the understanding that two people viewing the same item from different points in space may see different things. Children with ASD often struggle with VPT [4]; this can impact their ability to understand and respond to the perspectives of others. As a result, some social interactions may prove challenging for children with ASD. However, recent research has shown that humanoid social robots can help autistic children improve their VPT skills [5].

These children often experience anxiety and stress in social situations, which can be compounded by negative experiences with peers or other individuals in their community. By using humanoid social robots, children can practice their VPT in a safe and controlled environment, without fear of negative consequences or judgment. Caregivers (e.g. therapists, teachers, and parents) can build on the interest displayed by children with ASD towards the robots and use them as mediator tools, tailoring the interaction to the specific needs of the children at any given time [2], [6], [7]. In addition, social and educational robots can be programmed to provide feedback and support for children with ASD as they work on their VPT. For example, a robot might provide positive reinforcement and encouragement for successful attempts at VPT, or provide constructive feedback for areas that need improvement. This type of support and feedback can help build confidence and motivation in children with ASD and can provide a foundation for further improvement [8].

Therefore, the use of humanoid social and educational robots should be understood as a mediator tool for researchers and educators; they are to be used for improving VPT skills in children with ASD. By providing opportunities for social interaction and practice, giving feedback and support, and creating a non-threatening and non-judgmental environment, humanoid social robots can play an important role in helping children with autism [9].

This paper aims to investigate the effects of causal explanations provided by a humanoid social robot in improving the VPT skills of children with ASD in a set of interactive games. These games were designed taking into account the results obtained in a retrospective study [10] analysing earlier interactions between autistic children and the Kaspar robot that took place in former studies [9]. These former studies were used to identify scenarios that may be amenable to causal explanations [10] and inform the initial choice of educational games which were later adjusted using co-creation with teachers. Previous research was also used to formalise actual causality and implement it in an interface

Research supported by The UKRI Trustworthy Autonomous Systems (TAS) Hub (EP/V00784X/1) and the UKRI TAS Node on Verifiability (EP/V026801/2). We kindly acknowledge that the retrospective study relies on data gathered in the Horizon 2020 project BabyRobot (687831).

[1] Robotics Research Group, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK. Email: {m.sarda-gou, g.lakatos, p.holthaus, b.robins, s.moros, l.wood, f.amirabdollahian2}@herts.ac.uk

[2] Department of Informatics, King's College London, UK. Email: {hugo.araujo, mohammad.mousavi}@kcl.ac.uk

[3]Garston Manor School: Specialist School for Autism, Learning Difficulties and Speech and Language, UK. Email: head@garstonmanor.herts.sch.uk

Fig. 1. The humanoid social robot Kaspar

on the Kaspar robot [11]. The here presented study goes beyond these earlier attempts and combines them in a newly designed rigorous experiment, using a robot equipped with a causal explanation engine in an intervention in a school, and evaluating the resulting system's effect on the learning experience of the students with ASD.

In this research, Kaspar will play several interactive VPT games with the children and will provide a number of pre-programmed causal explanations through a remote-controlled interface. The alternative hypothesis of the study is that the causal explanations will have a positive impact on the children's VPT skills, reducing the mistakes and increasing the correct actions that the children will perform.

The remainder of this paper is organised as follows. In Section II, we provide a detailed account of our methodology, detailing the type of educational games used for our study, the type of causal explanations presented to children, and their validation in an interactive experiment. In Section III, we present the data obtained in this experiment and analyse them rigorously. In Section IV, we review the results and discuss whether causal explanations can improve learning of visual perspective for children with autism. We also present the directions of our future research before we conclude the paper in Section V.

II. METHOD

The study presented in this paper used the Kaspar robot (Figure 1), a social and educational companion robot developed by researchers at the University of Hertfordshire. It was specifically designed to help children with ASD develop social interaction and communication skills [12]. It is a child-size robot that has been purposefully designed with simplified, realistic human-like features offering a more predictable form of communication, making social interaction simpler, and more comfortable for the child. Kaspar has a child-like appearance and is approximately 56cm tall. The robot is equipped with sensors and cameras that allow it to respond to external stimuli, and it is capable of a range of movements, gestures and facial expressions (e.g. eye movements, blinking, nodding, shaking its head, waving its arms, opening its mouth and smiling, portraying 'happy' or 'sad' expressions, etc.). Kaspar is mainly used as an educational mediator in interaction with other people (peers and adult caregivers) and is particularly designed to engage children in activities and games that encourage them to interact with the robot and to develop their social skills (turn-taking, joint attention, cause and effect understanding, etc.). In addition, by using the robot as a tool for research and observation, researchers can gain new insights into the social and communication skills of children with ASD and develop new strategies for helping them to overcome these challenges.

A. Participants

Ten children with ASD took part in the study in their school. They have been selected based on the advice of their teachers who know them and could assess their suitability for participating in the study, for example, whether they have difficulties understanding VPT. This process resulted in nine male and one female participant in an age range between seven and ten years. They were divided randomly into the ECE or CEC groups to filter out interjection effects. This study has been approved by the University of Hertfordshire's ethics committee for studies involving human participants, protocol number: SPECS/SF/UH/04944. Informed consent was obtained in writing from all parents of the participating children.

B. Game design

The games were created taking into account the results from a former retrospective study [10]. The objective was to have four games related to VPT that elicited causal explanations so the robot had the chance to provide the children feedback to improve their VPT skills. As with previous studies involving Kaspar, its behaviour is under the control of a researcher so they can adjust to each child's specific needs and can make sure they have a good experience interacting with the robot [2]. In total, children were asked to respond to 24 questions that were asked by Kaspar. The following games were presented in order of ascending difficulty [3], ranging from out-of-sight positions and line-of-sight blockers, which are considered easier to understand, to the more difficult understanding of different perspectives on the same object:

1) Game 1: Show me the animal: Six laminated pictures of six different animals were placed around the room. Kaspar asked the child to show it some specific animal. In the experimental sessions, if the child placed the animal in front of Kaspar's eyes, they would receive positive feedback from the robot. If the children failed a trial by making a mistake like, for example, placing the picture too low, the robot would explicitly explain the reason why it could not see the animal until the children placed the picture in the right position. In the control sessions, children did not receive any feedback. This game had six trials.

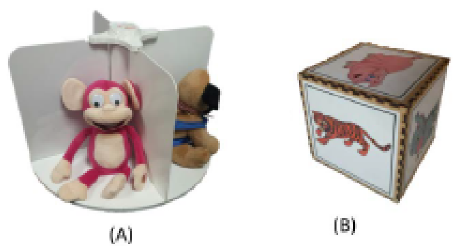

Fig. 2. (A) Turning table: the animal is placed on one side so only Kaspar or the child can see the toy. (B) Cube: there is a picture of an animal on each side. The child is asked to hold the cube so Kaspar can see a specific animal.

2) Game 2: Show me the animal on the cube: In this game, the children were holding a cube that had a picture of a different animal on each side, see Figure 2 (B). As in Game 1, Kaspar asked the child to show a specific animal to it. In the experimental sessions, if the child found the requested animal on the cube and turned the cube correctly so as to show the requested animal to Kaspar (whilst the child was seeing a different animal on the other side of the cube), they received positive feedback from the robot. If the children failed a trial, the robot explained the reason why it could not see the animal until the children placed the cube in the right position. In the control sessions, children did not receive any feedback. This game had 6 trials.

3) Game 3: Turning Kaspar's head: The researchers positioned 6 laminated pictures of 6 different animals around the room. Kaspar then asked the child to show it some specific animal. In this game, the child had to move Kaspar's head so the robot could see the animals. In the experimental sessions, if the child moved Kaspar's head correctly, they received positive feedback from the robot. If the children failed a trial, like, for example, moving the head too far left or too far right, the robot explicitly explained the reason why it could not see the animal until the children moved the head correctly. In the control sessions, children did not receive any feedback. This game had 6 trials.

4) Game 4: The turning table: This game involved the use of the turning table, cf. Figure 2 (A), and was divided into two parts, which were both repeated three times. One animal was placed on each side of the partitions on the table. Kaspar asked the child to show it a specific animal. The child then had to spin the turning table until the right animal was on Kaspar's side of the table so only Kaspar could see that specific animal. In the second part of the game, the researcher was the one spinning the table placing a specific animal in front of Kaspar. Then, the robot asked the child "What animal can I see?" to which the child had to say the specific animal that was in front of Kaspar on the turning table. In both parts of the game, in the experimental sessions, the children received positive feedback every time they performed a successful trial and they received a causal explanation when they failed a trial. In the control sessions, children did not receive any feedback.

$$\begin{aligned}
&\mathcal{F}_{canKasparSee}() = isKasparAwake = correct \wedge\\
&\quad areKasparsEyesClear = correct \wedge\\
&\quad isKasparsViewClear = correct\\
&\mathcal{F}_{canKasparSeeChosenAnimal}() = canKasparSee \wedge\\
&\quad chosenAnimal = correct \wedge\\
&\quad chosenAnimalPosition = correct \wedge\\
&\quad chosenAnimalRotation = correct \wedge\\
&\quad chosenAnimalHeight = correct \wedge\\
&\quad chosenAnimalDistance = correct
\end{aligned}$$

Fig. 3. Causal model: the variables and equations that model the system and represent the behaviour of the interaction between the children and Kaspar.

C. Conditions

The study was carried out following an ECE – CEC design (E = Experimental, C = Control) to look for potential effects of explanations on the effectiveness and persistence of each child's learning. The children were randomly assigned to one of the two groups and all of them had three sessions with the robot in a time frame of two weeks with at least one day in between sessions. The children who were in the ECE condition had first an experimental session followed by a control session and finally an experimental session again. The children in the CEC group did the opposite. In order to investigate how the causal explanations affected the improvement of the VPT skills, in the experimental sessions, the children received positive feedback every time they succeeded in a trial and a causal explanation when they failed a trial. In the control sessions, they only received positive feedback when they were successful but no casual explanations when they failed in the trials.

D. Causal model

A causal model was developed and implemented in Kaspar [11] so the robot could deliver a causal explanation every time the children failed a trial showing an animal to Kaspar. The causal model is based on Halpern-Pearl's theory of actual causality [13]. It comprises variables that represent the situations in the above-specified games and the observed interactions between the child and the robot. It also comprises equations that capture how different observed phenomena lead to different effects, such as Kaspar's view being obstructed or Kaspar viewing a different animal than intended. The model is instantiated on-the-fly according to what is taking place during a session and is fed into a causal analysis process that explores the model and infers the cause of the particular effect of interest automatically. A brief overview of the causal model is provided in Figure 3. In this figure, the equations and the variables that model the system are shown. The effect of interest with regard to whether Kaspar is physically able to see the animal (which is dictated by whether the view is not obstructed) and whether the animal is correctly shown (which is affected by whether the animal is in his field of view, correct rotation, and distance).

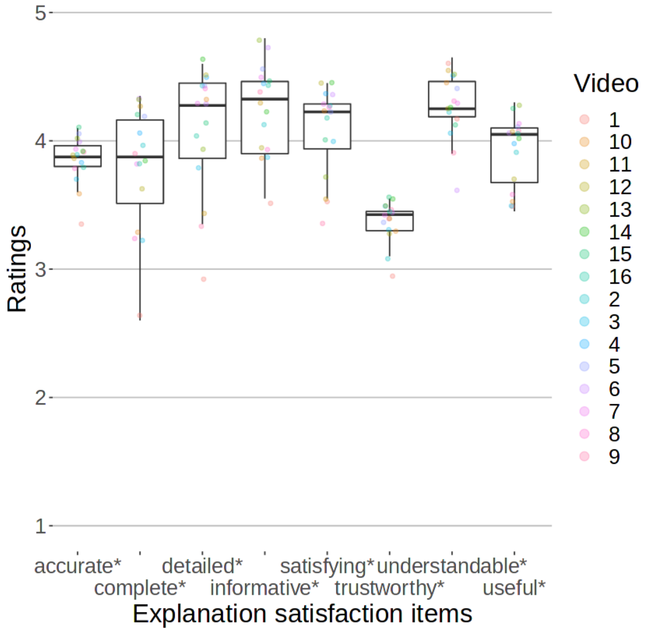

Fig. 4. Results of the evaluation of explanations in the multi-centre study. Results of the Explanation Satisfaction (ES) scale (5-point Likert scale) grouped by video number. Coloured points indicate the mean values of the other dimension. Asterisks mark items significantly greater than the average value.

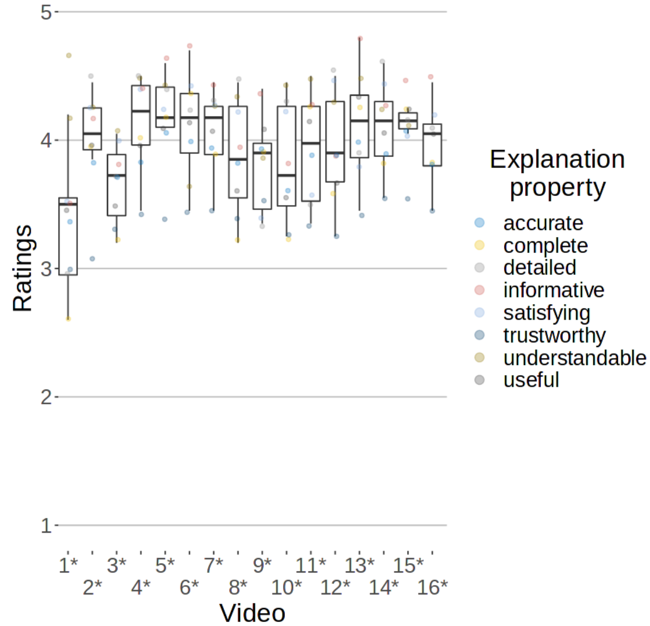

Fig. 5. Results of the evaluation of explanations in the multi-centre study. Results of the Explanation Satisfaction (ES) scale (5-point Likert scale) grouped by explanation property. Coloured points indicate the mean values of the other dimension. Asterisks mark items significantly greater than the average value.

The main aim is to make children understand the reason why the robot could not see the animal so they could rectify and show it correctly. As an example of the explanations provided by the causal analysis, if the child placed the cube too close to Kaspar's face, then the robot would say "*I cannot see the animal because it is too close to my eyes.*" or "*I cannot see the animal because my eyes are covered.*". The procedure was designed to follow a Wizard of Oz approach; while the children were playing, the robot's causal explanations and other reactions were triggered by the researchers, who were operating the robot remotely with a keyboard using the implemented causal model. We provide 16 different causal explanations [11], of which two are with respect to the distance between the animal and the robot, six are about the position of the animal (e.g. too far to the left), three explain that Kaspar's head was not moved correctly, further three state that the robot cannot see the animal because the vision is blocked, and two of the explanations are related to the child holding the wrong picture or the wrong object. Some of the explanations are specific to the type of game (e.g. moving Kaspar's head) but some could be used in all games.

III. Results

In this section, we first reflect on the results of our earlier verification of causal explanation validity to assess whether Kaspar presents generally understandable explanations to the children. Afterwards, we present the results of our interactive study to see whether these explanations have an effect on the children's learning about VPT.

A. *Causal explanations validation*

An initial survey has been carried out to assess the explanations generated by the system in terms of their general validity [11]. For this purpose, we asked 20 adult participants (10 research students or staff members from research groups based at King's College London and the same number from the University of Hertfordshire) to watch videos of Kaspar providing an explanation and then rate each explanation using the Explanation Satisfaction (ES) scale [14]. This survey is based on several key attributes of explanations such as whether they are understandable, satisfying, sufficiently detailed, complete, informative about the interaction, useful, accurate, and trustworthy. These attributes are used to assess the suitability of an explanation provided by an autonomous system. We additionally employed the Negative Attitude towards Robots Scale (NARS) to calibrate the obtained results against potential biases against robots. No other data, i.e. no personal data, was collected and the study was approved by the University of Hertfordshire's ethics committee for studies involving human participants, protocol number: SPECS/SF/UH/04944. In total, we showed 16 videos[1] to participants that contained all possible explanations for the variables of the causal network of the interactive games.

Because participant ratings were not normally distributed, we used the non-parametric one-sample Wilcoxon rank-sum test to test whether ratings on the ES scale were greater than the mean value. Results of the analysis [11], are summarised in Figure 4 (by ES item) and Figure 5 (by video sequence).

[1]A playlist of videos showing the interaction can be found on YouTube.

They attest that, when averaging across all the videos, each of the explanations is rated significantly above the neutral value (all $p < 0.001$). Likewise, ratings across the explanations are rated above neutral for each of the videos (all $p < 0.001$). Participant ratings on NARS attested a low negative attitude towards robots with mean values for $S1 \approx 1.78$ (interaction subscale), $S2 \approx 2.7$ (social subscale), and $S3 \approx 1.48$ (emotion subscale). S1 and S3 are rated significantly below the neutral value (both $p < 0.001$) whereas S2 could not be reliably distinguished from neutral ($p \approx 0.053$).

This confirms that, with neurotypical adults, the explanations that the system can generate are beneficial to relate cause and effect. Participants consistently rated them as accurate, complete, sufficiently detailed, satisfying, understandable, useful to their goals, and informative about the interaction. Knowing that adults find the generated explanations useful gives us an estimate of whether the generated explanations have the potential to help autistic children.

B. The impact of causal explanations

All child-robot interactions (N = 30) were video-recorded and coded in order to observe differences between control and intervention. The coding scheme included the correct answers after a trial, the causal explanations, the incorrect answers (both right after a trial and after an explanation), rectifications (both after an explanation or without explanation) and the total number of trials. We had these parameters for each game and each session of the trial. The videos were coded by a member of the research team and 20% of the videos were re-coded by a different member of the team. There was a strong agreement between the two raters ($\kappa = 0.93, p < .001$). Any disagreement was resolved through discussion.

The ratio of correct actions (RC) over the total number of actions (both correct and incorrect) was taken as a suitable parameter for our analysis. In order to obtain this value, we followed the equation in which, $c =$ the total number of correct actions, and $i =$ the total number of incorrect actions ($c + i$ equals the total number of sactions):

$$RC = \frac{c}{c + i}$$

A one-way ANOVA comparing the control and experimental sessions shows a significant difference between the ECE and CEC groups for RC ($F(1, 28) = 4.461, p = .04, \eta^2 = .14$). Ratio for incorrect actions over total number of actions was also calculated but given the complete opposite nature of such parameter, statistics obtained are redundant repeating the findings and hence not reported.

An independent sample t-test analysing only the first session and comparing the two conditions C and E reveals that there was a significant difference between the children who received causal explanations (who had a higher ratio of correct actions) and the children in the control session in their ratio of correct actions ($t(8) = -4.199, p = .003, 95\%$ CI, -0.43 to -0.13, Cohen's d $= 2.66$). Performing the same analysis to compare the two conditions in the second session resulted in ($t(8) = -.027, p = .979$) and again in the third session ($t(8) = -1.206, p = .262$), which indicates that after the first session, there were no more significant differences between the two groups. However, we can observe a reduction of the p-value in session 3, showing that the differences between groups increased after session 2. The impact of causal explanations is presented in Figure 6. Table I shows the comparison between sessions with and without explanations, where it can be also observed that the children were more successful in sessions with explanations.

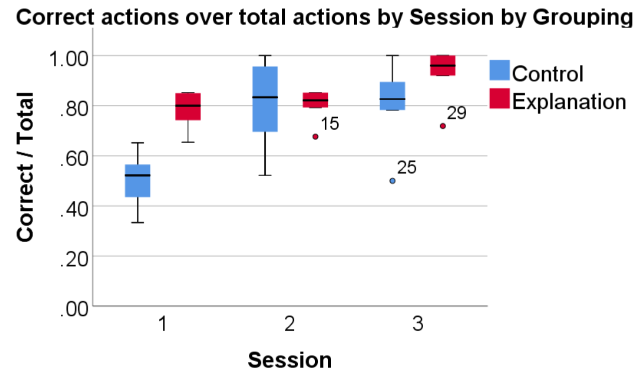

Fig. 6. Differences in the ratio of correct actions over total actions, observed for each session of the trial, for both control and intervention sessions. Children experience the sessions in order CEC (blue-red-blue) or ECE (red-blue-red).

TABLE I
CASE SUMMARIES COMPARING SESSIONS

	Control	Explanation
Correct	# 238 (70%)	# 323 (82%)
Mistakes	# 100 (30%)	# 71 (18%)
Total	# 338	# 394

IV. Discussion

The results obtained present convincing evidence in favour of the use of causal analysis and explanations in the trial sessions. Modelling the interaction between Kaspar and the children as a causal model and exploring the model to mathematically determine actual causality (according to Halpern-Pearl's theory) made it appropriate to accept the alternative hypothesis. A number of conclusions can be drawn from Figure 6 and Table I. For CEC: Control (blue in Figure 6) is followed by an explanation (red), which is then eventually followed by a control session (blue). Session 2 clearly shows an increased correctness ratio compared to session 1, which is then maintained in session 3, indicating a direct learning effect of adding explanations to Kaspar's behaviour. For ECE: The explanation session (red) is followed by a control session (blue), then followed by another explanation session (red). It is interesting to note that, while the average correctness ratio appears to be similar between all sessions, there seems to be an increase between the first and the third session. However, in session 2 without explanations, there is a higher variance, which might be explained by children relying on these expectations.

In total, the ECE group has better results in the count of correct answers and thus fewer mistakes compared to the CEC group. This could be explained by the fact that the former group had more sessions with explanations than the latter. Based on the results of this study, it can be concluded that using a social humanoid robot to provide causal explanations can be an effective tool in improving VPT in autistic children. Specifically, the effect of condition in the first session suggests that this approach was most effective the first time the children interacted with the robot providing causal explanations, which may indicate that the provision of feedback by the robot played a key role in the improvement of VPT. That is to say, the results indicate that the explanations had an initial positive effect on the children's VPT. The data shows the biggest difference in session 1, indicating that once the children had been exposed to the explanations, they levelled up their skills and retained them until the last session. These further results suggest that causal explanations are a good way to improve the VPT in children with ASD since the children understood the causal explanations given by the robot, and applied this feedback to their VPT and preserved this knowledge.

The findings of this study have important implications for the design and implementation of interventions aimed at improving VPT in autistic children. The use of robots in this context can provide a more engaging and interactive experience for children, which could lead to better outcomes [5]. Following the results of the here presented study, researchers and practitioners may want to consider commonly using causal explanations when using robots for improving VPT in autistic children.

However, it should be noted that this study has some limitations. First, the sample size was relatively small, which limits the generalizability of the results. Second, the study only investigated the short-term effects of the robot intervention, and it is unclear whether the observed improvements in VPT would persist over longer periods of time. Therefore, future studies with larger sample sizes and longer follow-up periods are needed to further investigate the efficacy of robot-explained VPT interventions.

V. Conclusion

In this paper, we have presented an investigation into the effects of causal explanations provided by an explanation engine using the Kaspar robot to improve VPT skills of children with autism in a set of interactive games. The findings of our study suggest that pre-assessed relevant causal explanations can be an effective tool for improving children's understanding of VPT. The results also support the potential of using a social robot as a tool for supporting social communication and interaction in autistic children. However, further research is needed to better understand the full potential of robot-assisted interventions for improving VPT in autistic children and to explore the optimal conditions for using robots in this context.

References

[1] L. Wing, The Autistics Spectrum: Aguide for parents and professionals. Constable and Company Limited, 1996.

[2] B. Robins and K. Dautenhahn, "Kaspar, the social robot and ways it may help children with autism–an overview," Enfance, vol. 2018, no. 1, pp. 91–102, 2018.

[3] A. F. d. C. Hamilton, R. Brindley, and U. Frith, "Visual perspective taking impairment in children with autistic spectrum disorder," Cognition, vol. 113, no. 1, pp. 37–44, 2009.

[4] N. Yirmiya, M. Sigman, and D. Zacks, "Perceptual perspective-taking and seriation abilities in high-functioning children with autism," Development and psychopathology, vol. 6, no. 2, pp. 263–272, 1994.

[5] G. Lakatos, L. J. Wood, D. S. Syrdal, B. Robins, A. Zaraki, and K. Dautenhahn, "Robot-mediated intervention can assist children with autism to develop visual perspective taking skills," Paladyn, Journal of Behavioral Robotics, vol. 12, no. 1, pp. 87–101, 2020.

[6] C. Huijnen, M. Lexis, and L. De Witte, "Robots as new tools in therapy and education for children with autism," International Journal of Neurorehabilitation, vol. 4, no. 4, pp. 1–4, 2017.

[7] L. J. Wood, A. Zaraki, M. L. Walters, O. Novanda, B. Robins, and K. Dautenhahn, "The iterative development of the humanoid robot kaspar: An assistive robot for children with autism," in Social Robotics: 9th International Conference, ICSR 2017, Tsukuba, Japan, November 22-24, 2017, Proceedings 9. Springer, 2017, pp. 53–63.

[8] J. Wainer, B. Robins, F. Amirabdollahian, and K. Dautenhahn, "Using the humanoid robot kaspar to autonomously play triadic games and facilitate collaborative play among children with autism," IEEE Transactions on Autonomous Mental Development, vol. 6, no. 3, pp. 183–199, 2014.

[9] T. Zorcec, B. Robins, and K. Dautenhahn, "Getting engaged: assisted play with a humanoid robot kaspar for children with severe autism," in ICT Innovations 2018. Engineering and Life Sciences: 10th International Conference, ICT Innovations 2018, Ohrid, Macedonia, September 17–19, 2018, Proceedings 10. Springer, 2018, pp. 198–207.

[10] M. Sarda Gou, G. Lakatos, P. Holthaus, L. J. Wood, M. R. Mousavi, B. Robins, and F. Amirabdollahian, "Towards understanding causality - a retrospective study of scenarios that trigger explanations in interactions between a humanoid robot and autistic children," in International Conference on Robot and Human Interactive Communication (RO-MAN 2022). Naples, Italy: IEEE, 2022.

[11] H. Araujo, P. Holthaus, M. Sarda Gou, G. Lakatos, G. Galizia, L. Wood, B. Robins, M. R. Mousavi, and F. Amirabdollahian, "Kaspar Causally Explains: Causal Explanation in an Assistive Robot for Children with Autism Spectrum Disorder," in International Conference on Social Robotics (ICSR 2022), ser. Lecture Notes in Computer Science, F. Cavallo, J.-J. Cabibihan, L. Fiorini, A. Sorrentino, H. He, X. Liu, Y. Matsumoto, and S. S. Ge, Eds. Florence, Italy: Springer Cham, 2022, vol. 13818, pp. 85–99.

[12] L. J. Wood, A. Zaraki, B. Robins, and K. Dautenhahn, "Developing kaspar: a humanoid robot for children with autism," International Journal of Social Robotics, vol. 13, pp. 491–508, 2021.

[13] J. Y. Halpern, Actual causality. MiT Press, 2016.

[14] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman, "Metrics for explainable ai: Challenges and prospects," arXiv preprint arXiv:1812.04608, 2018.