To appear: 2019 IEEE-RAS-EMBS International Conference on Rehabilitation Robotics (ICORR), 24-28 June, 2019, Toronto, Canada, 6 pages.

Supporting Play by Applying Haptic Guidance Along a Surface Learnt from Single Motion Trajectories

Javier L. Castellanos-Cruz[1], Mahdi Tavakoli[2], Patrick M. Pilarski[3], and Kim Adams[1]

Abstract— **Haptic-enabled teleoperated robots can help children with physical disabilities to reach toys by applying haptic guidance towards their toys, thus compensating for their limitations in reaching and manipulating objects. In this article we preliminarily tested a learning from demonstration (LfD) approach, where a robotic system learnt the surface that best approximated to all motion trajectories demonstrated by the participants while playing a whack-a-mole game. The end-goal of the system is for therapists or parents to demonstrate to it how to play a game, and then be used by children with physical disabilities. In this study, four adults without disabilities participated, to identify aspects that will be necessary to improve before conducting trials with children. During the demonstration phase, participants played the game in normal teleoperation, assuming the role of the therapist/parent. Then, the surface was modeled using a neural network. Participants played the game without and with the haptic guidance. The movements of the robotic system were mirrored to induce errors in movements, and thus require the guidance. Participants spent more time, moved the robot longer distances, and had jerkier movements when they played the game with the guidance than without it. Possible reasons were discussed, and several solutions were proposed to improve the system. The main contribution of this paper was the learning of a surface instead of learning a single motion trajectory.**

I. INTRODUCTION

Play contributes greatly to child development, especially for the development of sensory, motor, cognitive, communication and social skills [1]. Play is, in essence exploratory, meaning that play activities involve movement, manipulation, and interaction with the environment [2]. However, children with physical disabilities may not have the physical capabilities to do so.

Children with physical disabilities, such as those resulting from cerebral palsy, experience limitations in play because of difficulties in reaching and handling objects, abilities that are required for manual tasks in everyday life or for playing [3]. Manipulative exploration of the environment is as important as visual exploration, since the physical interaction can provide additional information about the objects that cannot be sensed visually, such as object properties like texture, weight, rigidity, or temperature [4]. When the ability of children with physical disabilities to play is affected, the direction of play is often led by others, thus, children with physical disabilities may miss opportunities to learn by playing [5].

[1]J. L. Castellanos-Cruz, and K. Adams are with Faculty of Rehabilitation Medicine, University of Alberta, Canada (e-mail: javierle@ualberta.ca and kdadams@ualberta.ca; phone: +1 780 492 5422).

[2]M. Tavakoli is with the Electrical and Computer Engineering Department, University of Alberta, Canada (e-mail: mahdi.tavakoli@ualberta.ca).

[3]P. M. Pilarski is with the Department of Medicine, University of Alberta, Canada (e-mail: pilarski@ualberta.ca).

Haptic-enabled robots can allow users to have the sense of touch. They can allow children to sense the mechanical characteristics of objects they interact with using the robot. Becerra et al. [6] tested a haptic-enabled teleoperated robotic system for sensing size, roughness, hardness, and shape of objects. Children were able to identify the characteristic of the object without looking at the objects.

Haptic-enabled robots can also apply haptic guidance for helping people with disabilities do manual tasks [7]. Haptic guidance has been employed to help children with cerebral palsy to improve Chinese handwriting [8]. Children wrote Chinese characters on a computer screen using a pen-like robot to follow the templates given as guidelines. If the child's handwriting was off the template, force feedback was provided to pull the child's hand towards and along the trajectory of the character. Haptic guidance has been beneficial for maneuvering an electric powered wheelchair [9], [10]. Children control the speed and direction of the wheelchair with force-feedback joystick, and the joystick exerts forces to direct the wheelchair in the correct direction if the wheelchair goes off the predefined path or if it goes in the direction of an obstacle within the activity/task.

Haptic guidance can contribute to better performance in a task, although, it may not necessarily be in accordance to the movement intentions of the users, thus, it may oppose to how users want to move [7], [11].

One of the issues with the aforementioned approaches is that the haptic guidance model (e.g., trajectory or path) is manually programmed. The problem is that robotic systems like those may not perform other tasks, e.g., write English characters, unless the haptic guidance models are reprogrammed. In the context of robots for play, the end-goal of most robotic systems is to have them at children's home so that children can play on a daily basis. However, it is likely that most parents will not have the necessary programming skills.

Robots can be trained to do tasks from demonstrations performed by humans. Robot learning from demonstration (LfD) is the field of study that explores how robots can learn from examples provided by a human that acts as a teacher [12]. For example, a helicopter was trained to fly autonomously by demonstrations from an expert pilot [13]. The robot learned the ideal trajectory from the demonstration by the pilot. Similarly, a robotic arm was trained to play mini golf from demonstrations by a golfer [14].

Learning low-level motion trajectories is a common strategy for teaching robots to do specific tasks [12], [15], and applied in rehabilitation robotics. People with upper limb amputations could teach their prosthesis how to move by showing desired movement trajectories with their non-amputated limb [16]. People with physical disabilities, such as stroke patients, could improve their motor functionality through repetitive motion trajectories that a therapist can teach to a haptic-enabled robot [17]. Thus, it follows that children with physical disabilities could use a haptic-enabled teleoperation system to reach their toys by applying haptic guidance along a demonstrated trajectory [18].

In the context of play, it would be necessary to learn all the motion trajectories to reach each toy in the environment. One issue with this is how a child could switch the haptic guidance to different trajectories to reach the different toys. An additional input (e.g., a button) could work, but it is not always feasible for a child to be able to control it, physically, or cognitively. We propose to apply a LfD approach to learn the surface that approximates the movements necessary to reach the toys in the environment and apply haptic guidance along the surface. A therapist or parent could demonstrate to the robotic system the movements to play a game with several targets, and then a haptic-enabled robotic system could support children with physical disabilities to play by applying haptic guidance.

Neural networks create models that map the inputs to outputs [19]. A neural network can create a model that approximates to a target continuous function, which is demonstrated by a set of examples (inputs). Thus, a neural work could be used to learn the function that best approximates to the movements demonstrated by a therapist or parent while playing a game.

The objective of this article is to test the concept of surface modelling as a LfD approach to support play in children with physical disabilities. A neural network was used to model the surface of the movements performed by adults without physical disabilities while playing a whack-a-mole game using a haptic telerobotic system. At this stage, the haptic telerobotic system was tested by adults without disabilities to ensure the system is safe to use for children, and to identify feasibility of this LfD approach and possible improvements necessary before conducting trials with children with physical disabilities.

II. METHODS

A. Study design

A cross-over experimental design was performed for testing the robotic system. Participants played the game without and with haptic guidance. Half of the participants started playing with guidance and the other half without guidance, and then they switched conditions. This way participants served as their own control.

B. Participants

Four adult university students without physical or cognitive disabilities participated in this study. The participants' age ranged from 23 to 44 years old (Mean=30.5, SD=9.6). Participants Ethical approval was obtained from the Health Research Ethics Board at the University of Alberta.

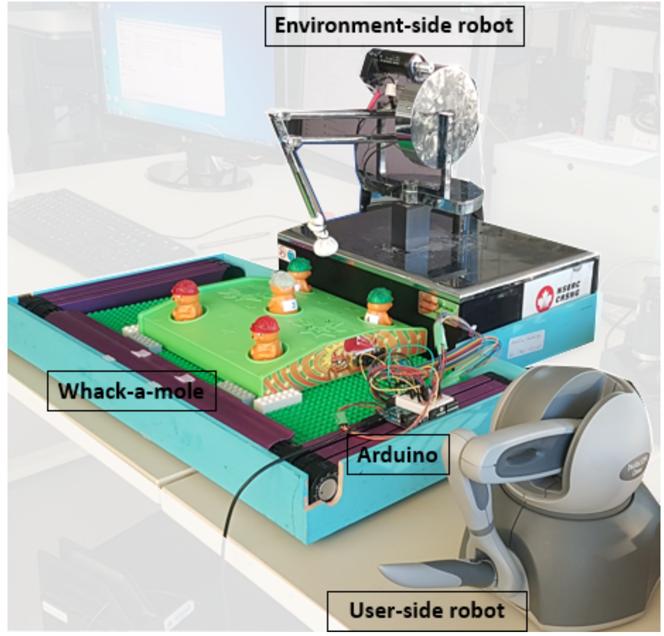

Fig. 1. Setup of the robotic system and the game.

C. Materials

1) Robotic system: The telerobotic haptic system included one Phantom Premium 1.5A and one Touch haptic robot (3D Systems, Inc., Rock Hill, SC, USA). The Phantom Premium was placed in the environment (slave side) where it interacted with the game. The other haptic robot was placed in the user side (master side), for participants to control. The robots were programmed in bilateral teleoperation mode using a PID controller based on position. The robots were programmed in Simulink R2017a Matlab/Simulink and used the Quarc V2.2 library (Quanser Inc., Markham, ON, Canada) on a Windows PC.

A whack-A-Mole Arcade Game by Fischer-Price was adapted to light up and turn off the lights of the five moles, and sense the pressing of the moles with switches, using an Arduino Leonardo microcontroller. The Arduino was in serial communication with the Robots' PC, sending the information about the moles that were pressed, in this way the data could be analyzed mole by mole. Fig. 1 shows the main components and setup of the robotic system and the game.

2) Surface modelling using a Neural Network: A neural network was trained to create a model of the 3D surface that best approximates to the movements demonstrated by the participants. A feed-forward neural network was used to learn the height (position in the Z-axis) with respect to the X and Y position of the environment-side robot. The neural network created a model as:

$$\hat{Z} = f(X, Y) \tag{1}$$

where $\hat{Z}$ is the output of the neural network. Fig. 3 illustrates the structure of the network implemented that created such a model. The structure was 2-9-1. The network

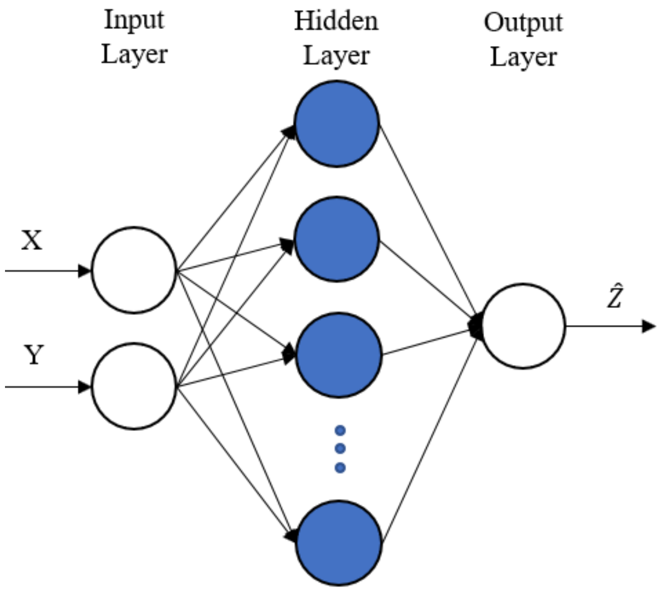

Fig. 2. Structure of the feed-forward neural network. Two input nodes, one hidden layer with nine hidden nodes, and a single output node.

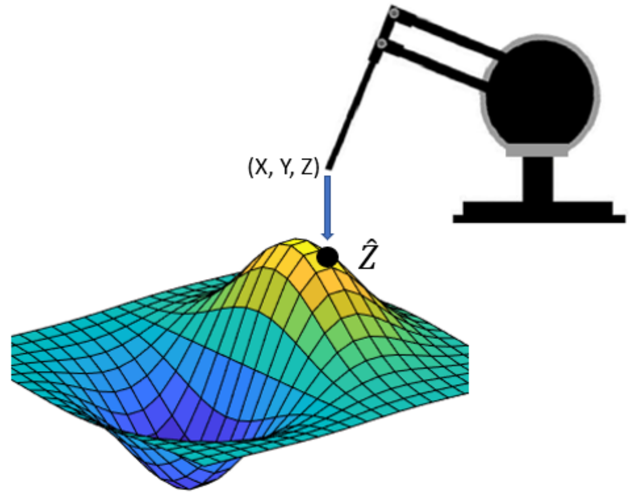

Fig. 3. Illustration of the haptic guidance.

had two inputs only: X and Y positions of the environment-side robot. The network had one hidden layer with nine hidden nodes. The number of hidden nodes was selected by visually exploring the surfaces generated having from 1 to 10 hidden nodes. Finally, the output layer had only one node, the approximate Z position of the environment-side robot.

3) Haptic guidance: Haptic guidance was applied towards the predicted height ($\hat{Z}$) by the neural network. A simple artificial potential field was implemented towards X, Y, $\hat{Z}$, where X and Y were the coordinates of current position of the environment-side robot. Fig. 3 illustrates the haptic guidance along a surface. Forces were applied in the Z-axis (i.e., upwards or downwards) of the user-side robot as:

$$F_Z = K(\hat{Z} - Z) = 100(\hat{Z} - Z) \tag{2}$$

The force was proportional to the difference between the true z-position and the predicted. K was set to 100 N/m.

D. Procedure

In this study, participants came to one session that lasted approximately half hour. Participants played the game three times using their non-dominant hand. In the demonstration phase they played the game in normal teleoperation. In the testing phase they played the game under two teleoperation conditions: mirrored teleoperation without guidance, and mirrored teleoperation with guidance. Normal teleoperation refers to the user-side and environment-side robots following each other's position. Mirrored teleoperation, in this study, refers to mirroring the Y- and Z-axes, thus when the user-side robot was moved upwards and towards the right, the environment-side robot moved downwards and towards the left. Mirroring of the teleoperation was done so the participants might move in the wrong direction, and thus need to use the haptic guidance to move in the correct direction. Without mirroring, the participants, who did not have motor impairments, may not have engaged the guidance feature.

Participants played the game under each condition until whacking 60 moles in total. The moles were randomly lit up in sets of three, after whacking those three, another three moles were lit up.

All participants started playing the game under normal teleoperation, the demonstration phase. Then, participants waited between five and ten minutes until the neural network was trained using the data they had demonstrated. In this case, the participants themselves demonstrated to the system how to play the game, but if the user was a child with physical disabilities, a therapist or parent could perform this demonstration.

After the neural network was trained, participants played the game under the conditions of mirrored teleoperation with and without haptic guidance. Two participants started playing with guidance and two participants without guidance, and then switched conditions. Participants waited about five minutes before switching conditions.

E. Data collection and analysis

While each participant played the game in the normal teleoperation condition during demonstration, the X, Y, and Z position of the environment-side robot was collected, at a sampling frequency of 200Hz. A dataset for each participant was composed, consisting of 60 episodes (i.e., the interval that participants took to whack each mole). The datasets were divided into training, test and validations sets, with 80%, 10%, and 10% of the data, respectively. The neural network was trained using the Levenberg-Manquardt back-propagation algorithm, and using mean squared error (MSE) as the performance, or cost, function. To avoid overfitting, training stopped when 100 epochs were reached, a minimum performance gradient of 5e-6 was reached, or the MSE reached $0.0001m^2$. The performance of the neural networks was measured using the MSE.

From all three teleoperation conditions the environment-side robot's position was recorded. Fourteen out of 720 episodes were excluded because of a malfunction with the game, i.e., a switch got stuck sometimes. From the included episodes the average time the user took to whack each mole, the distance travelled, and the jerkiness of the movements, were measured. The measure of jerkiness can reflect how the participants' movements were affected by mirroring the teleoperation and how the haptic guidance supported the participants to complete the activity. Jerkiness was measured by using the Dimensionless Jerk formula:

$$LDLJ = -log\left(\frac{(t_2 - t_1)^3}{{v_{peak}}^2} * \int_{t1}^{t2} |\frac{d^2v(t)}{dt^2}|^2 \, dt\right) \quad (3)$$

where v is the velocity at which the environment-side robot was moving. LDLJ is a valid measure for measuring smoothness of movements [20]. The lower the value of LDLJ the jerkier the movements.

Linear mixed-effects models were used to statistically compare the results between conditions. The results of time, distance, and jerkiness of each participant in each teleoperation condition were compared. The significance level of the statistical tests was 0.05. At the end of the study, participants were asked which condition was the easiest and the hardest for playing the game, and why. Their responses were recorded into the research notes.

III. Results

The performance in terms of MSE of the neural networks that were trained for each participant were: 0.000189, 0.000203, 0.000151, and 0.000265m^2. Fig. 4A illustrates a few randomly chosen motion trajectories demonstrated by a participant while he was playing in the normal teleoperation condition, and Fig. 4B illustrates the surface that was created by the neural network that obtained the highest MSE (2.0368e-04m^2).

Table I lists the means and standard deviations for time, distance, and jerkiness (LDLJ), for each participant when they played under the conditions of normal teleoperation, mirrored teleoperation without guidance, and mirrored teleoperation with guidance. Additionally, it lists the p-values of the linear mixed-effects model for comparisons between normal teleoperation and mirrored teleoperation without guidance, and mirrored teleoperation without guidance and mirrored teleoperation with guidance.

After participants played the game under the three teleoperation conditions, all of them reported that the easiest condition in which to play the game was when the system was in normal teleoperation, and the hardest was mirrored teleoperation with guidance. They commented that the surface created by the neural network had a shape that was not aligned with their movements. Also, they commented that the moles were harder to whack when the system had the guidance on.

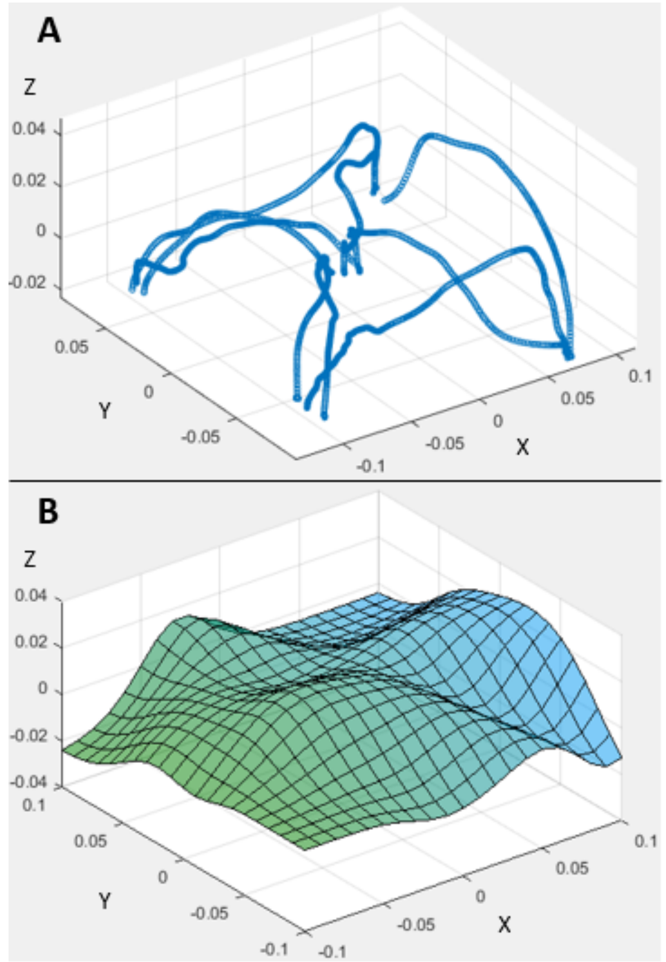

Fig. 4. Surface created by a neural network. A. Movements demonstrated by a participant while playing in normal teleoperation. B. The surface created by a neural network from the movements demonstrated in A.

IV. Discussion

Neural networks can be employed to model the surface in which the robot moves as seen in Fig. 4. A therapist or a parent could demonstrate the movements the robot should do to complete an activity. In this study, the participants assumed that role by demonstrating to the system how they, themselves, would play the game when the system was in the normal teleoperation condition.

Visual examination of the movements of each participant during demonstration revealed that there was a high variability in the movements performed by the participants. In this study, each participant demonstrated the movements to whack 60 moles, but perhaps fewer movement demonstrations are necessary and could result in smoother surfaces. The neural network was trained to minimize the MSE, thus, in some sense, the neural network learnt the average height (position in the Z-axis) with respect to the X and Y positions demonstrated and tried to generalize to the X and Y positions that were not demonstrated. Some surfaces had steep slopes, perhaps due to overfitting. A possible solution is to regularize (i.e., controlling the complexity of the surface models created

TABLE I

COMPARISON BETWEEN NORMAL TELEOPERATION AND MIRRORED TELEOPERATION WITHOUT GUIDANCE

	Normal teleoperation	Mirrored teleoperation without guidance	Mirrored teleoperation with guidance	p-value Normal vs. Mirrored without guidance	p-value Mirrored without vs. with guidance
Time (s)	1.132 ± 0.661	2.806 ± 2.404	3.903 ± 3.793	0.000	0.000
Distance (m)	0.221 ± 0.106	0.353 ± 0.275	0.459 ± 0.367	0.000	0.000
Jerkiness	-9.165 ± 0.935	-10.188 ± 1.349	-10.729 ± 1.374	0.000	0.000

by using a penalty factor in the optimization algorithm) the training of the neural network, this way overfitting could be diminished. Also, a different cost function could be created, for instance, a cost function that takes into account the slope of the surface, in addition to the MSE. Additionally, the surface modelling could be improved if the therapist or parent makes movements over the entire play area. In this activity, there were empty spaces where the participants did not move the robot, as seen in Fig. 4A.

All participants commented that the haptic guidance applied along the surface was not aligned with their movements. A possible reason that explains the misalignment with the participants' movements is that the neural network was trained with ambiguous data. The neural networks were trained with movements that overlapped each other but had different targets. For example, trajectories were generated when the participants moved from the left-bottom mole to the mole in the middle, and when they moved from the left-bottom mole to the right-upper mole. The neural network learnt and modeled the surface as mixture of both movements, thus, making it difficult to attain the motion trajectories that could reach both moles with low error. For activities like this whack-a-mole game, it will be necessary to create different surface models that guide the user to each mole and activate the surface that guides the user from the current mole to the target mole.

Mirroring the axes of the teleoperation certainly increased the difficulty for completing the activity. According to Table I, when the system was in mirrored teleoperation without guidance, participants spent significantly more time, moved the robot longer distances, and had jerkier movements to whack each mole than when they played in normal teleoperation. Mirroring the axes of the teleoperation caused participants to move involuntarily in the wrong directions and not efficiently. To some extent, mirroring of the axes simulated the movements that a person with physical disabilities might perform, especially if he/she experiences involuntary movements.

Haptic guidance did not help the participants to complete the activity. According to Table I, participants spent significantly more time, moved the robot longer distances, and their movements were jerkier when they played the game in the mirrored teleoperation with guidance condition than the without guidance condition. One reason was that the surface model created by the neural network was above the moles, even at the X and Y positions of the moles, and when participants wanted to push straight down on the moles, the haptic guidance along the surface was against those movements. Therefore, participants had to push harder to overcome the force of the haptic guidance and be able to whack the moles. In this study, the K force constant was set to 100N/m, if the constant had been increased it to higher value, participants may not have been able to whack the moles. A possible solution is to switch the type of haptic guidance when the environment-side robot's end-effector is close to a mole, so that the users can push straight down easier.

The haptic guidance applied along the surface created by the neural network should have helped the participants to not have to think about controlling the robot in the Z-axis. Thus, participants should have been able to devote most of the cognitive effort to control the robot in the two remaining degrees of freedom, the X and Y axes. This feature could be beneficial for children with physical disabilities. However, to accomplish this it will be necessary to improve the surface modelling so that the haptic guidance is aligned with the users' movements.

There were limitations in this study. There were only four participants, however, the large number of episodes (number of moles whacked) allowed statistical analysis. Only adults without physical impairments tested the system, and not children with physical disabilities, who are the target population. However, the findings of this study were helpful to identify the possible improvements that are necessary before doing trials with children. In addition, there was a washout period of about five minutes between the conditions of mirrored teleoperation without and with guidance, but results could have been different with a longer washout period, although, counterbalancing was implemented to help control for learning effects.

This work showed that modelling of a surface is possible, but future research is needed. We will improve the surface modelling by implementing the suggestions above: regularizing the learning, creating a new cost function, or training the neural network with movement examples that cover the entire play area. Additionally, it will be necessary to train the neural network with data that it is not ambiguous or perhaps create multiple surface models. Haptic guidance will be also implemented to help participants push straight down on the moles. Other tasks can be trialed to identify which activities are best for this LfD approach and which could be more useful for children with physical disabilities. The guidance along a surface will be compared to other guidance methods such as guidance on single motion trajectories, artificial potential fields, and forbidden region virtual fixtures. Finally, we will test the system with children with and without physical disabilities after the improvements are done.

V. Conclusions

This study has introduced a learning from demonstration approach based on surface modelling. A neural network was used to model a surface that represented the movements that adults performed while playing whack-a-mole using a haptic telerobotic system. The main contribution of this approach was the fact that it is possible to learn, not just a single motion trajectory as commonly done in LfD, but the surface that best represents multiple motion trajectories. This approach could benefit children with physical disabilities, since it could help children to play and reach their toys or objects since they only have to worry about moving the robot in the X and Y directions, and the surface will aid them in the Z direction. However, the surface modelling needs some improvements before conducting trials with children, so that it can better match the movements demonstrated by the therapist or parent.

Acknowledgment

This research was funded by Collaborative Health Research Projects (CHRP), a joint initiative of the National Sciences and Engineering Research Council (NSERC) and Canadian Institutes of Health Research (CIHR). Grant numbers: 462227-14 and 134744.

References

[1] K. J. Tanta and S. H. Knox, "Play," in *Occupational Therapy for Children and Adolescents*, 7th ed., J. Case-smith and J. C. O'Brien, Eds. St. Louis, USA: Elsevier, 2014, pp. 483–495

[2] K. Stagnitti and C. Unsworth, "The importance of pretend play in child development: An occupational therapy perspective," *Br. J. Occup. Ther.*, vol. 63, no. 3, pp. 121–127, 2000.

[3] A. C. Eliasson et al., "The Manual Ability Classification System (MACS) for children with cerebral palsy: scale development and evidence of validity and reliability," *Dev. Med. Child Neurol.*, vol. 48, no. 07, pp. 549–554, 2006.

[4] L. Fenson and R. Schell, "The origins of exploratory play," *Early Child Dev. Care*, vol. 19, no. 1–2, pp. 3–24, 1985.

[5] A. M. Okimoto, A. Bundy, and J. Hanzlik, "Playfulness in children with and without disability: Measurement and intervention," *Am. J. Occup. Ther.*, vol. 54, pp. 73–82, 2000.

[6] L. Becerra, T. Antunes, H. M. Capel, A. Wiebe, and K. D. Adams, "Testing of an assistive robot system for haptic exploration of objects," *Assist. Technol.*, pp. 1–9, 2018.

[7] N. Jafari, K. Adams, and M. Tavakoli, "Haptics to improve task performance in people with disabilities: A review of previous studies and a guide to future research with children with disabilities," *J. Rehabil. Assist. Technol. Eng.*, vol. 3, pp. 1–13, 2016.

[8] K. S. Choi and K. H. Lo, "A hand rehabilitation system with force feedback for children with cerebral palsy: two case studies," *Disabil. Rehabil.*, vol. 33, no. 17–18, pp. 1704–1714, 2011.

[9] L. Marchal-Crespo, J. Furumasu, and D. J. Reinkensmeyer, "A robotic wheelchair trainer: design overview and a feasibility study," *J. Neuroeng. Rehabil.*, vol. 7, no. 40, pp. 1–12, Jan. 2010.

[10] X. Chen, C. Ragonesi, J. C. Galloway, and S. K. Agrawal, "Training Toddlers Seated on Mobile Robots to Drive Indoors Amidst Obstacles," *IEEE Trans. Neural Syst. Rahabilitation Eng.*, vol. 19, no. 3, pp. 271–279, 2011.

[11] I. Sakamaki et al., "Preliminary testing by adults of a haptics-assisted robot platform designed for children with physical impairments to access play," *Assist. Technol.*, 2017.

[12] S. Chernova and A. L. Thomaz, *Robot Learning from Human Teachers*. Morgan and Claypool Publishers, 2014.

[13] A. Coates, P. Abbeel, and A. Y. Ng, "Learning for Control from Multiple Demonstrations," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 144–151.

[14] S. M. Khansari-Zadeh, K. Kronander, and A. Billard, "Learning to Play Minigolf: A Dynamical System-based Approach," *Adv. Robot.*, vol. 26, no. 17, pp. 1967-1993, 2012.

[15] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Rob. Auton. Syst.*, vol. 57, no. 5, pp. 469–483, 2009.

[16] G. Vasan and P. M. Pilarski, "Learning from Demonstration: Teaching a Myoelectric Prosthesis with an Intact Limb via Reinforcement Learning," in *2017 International Conference on Rehabilitation Robotics (ICORR)*, 2017, pp. 1457–1464.

[17] C. Martinez and M. Tavakoli, "Learning and Robotic Imitation of Therapist's Motion and Force for Post-disability Rehabilitation," in *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2017, pp. 2225–2230.

[18] M. Najafi, M. Sharifi, K. Adams, and M. Tavakoli, "Robotic assistance for children with cerebral palsy based on learning from telecooperative demonstration," *Int. J. Intell. Robot. Appl.*, vol. 1, no. 1, pp. 43–54, 2017.

[19] S. Haykin, "Introduction," in *Neural Networks and Learning Machines*, 3rd ed., vol. 3, New York: Pearson, 2008, pp. 1–46.

[20] S. Balasubramanian, A. Melendez-Calderon, A. Roby-Brami, and E. Burdet, "On the analysis of movement smoothness," *J. Neuroeng. Rehabil.*, vol. 12, no. 1, pp. 1–11, 2015.