



Adaptation of the Difficulty Level in an Infant-Robot Movement Contingency Study

José Carlos Pulido¹(✉), Rebecca Funke², Javier García¹, Beth A. Smith²,
and Maja Matarić²

¹ Universidad Carlos III de Madrid, Madrid, Spain
{jcpulido,fjgpolo}@inf.uc3m.es

² University of Southern California, Los Angeles, USA
{rfunke,beth.smith,mataric}@usc.edu

Abstract. This paper presents a personalized contingency feedback adaptation system that aims to encourage infants aged 6 to 8 months to gradually increase the peak acceleration of their leg movements. The ultimate challenge is to determine if a socially assistive humanoid robot can guide infant learning using contingent rewards, where the reward threshold is personalized for each infant using a reinforcement learning algorithm. The model learned from the data captured by wearable inertial sensors measuring infant leg movement accelerations in an earlier study. Each infant generated a unique model that determined the behavior of the robot. The presented results were obtained from the distributions of the participants' acceleration peaks and demonstrate that the resulting model is sensitive to the degree of differentiation among the participants; each participant (infant) should have his/her own learned policy.

Keywords: Socially assistive robotics · Infant-robot interaction
User adaptation · Reinforcement learning

1 Introduction

Infants produce a variety of movements in order to modulate task-specific actions such as reaching, crawling, and walking [1, 2]. Through a dynamic process of exploration and discovery, they learn how to control their bodies and interact with their environments. In contrast to typically developing (TD) infants, infants at risk (AR) for developmental delays often have neuromotor impairments

This work was supported by NSF award 1706964 (PI: Smith, Co-PI: Matarić). In addition, this work was developed during an international mobility program at the University of Southern California being also partially funded by the European Union ECHORD++ project (FP7-ICT-601116), the LifeBots project (TIN2015-65686-C5) and THERAPIST project (TIN2012-38079).

involving strength, proprioception, and coordination. These challenges can lead to greater difficulty with movement and potentially a decreased motivation to move and explore.

Past works have used wearable sensors and/or 3-dimensional motion analysis systems to assess differences in movement patterns between infants with TD and infants AR or with developmental delays. Studies have demonstrated that movement variables such as kicking frequency, spatiotemporal organization, and interjoint and interlimb coordination are different between infants with TD and infants AR [3], with intellectual disability [4], with myelomeningocele [5,6], with Down syndrome [7], or born preterm [8]. Studies have also shown that the acquisition of new motor skills is correlated to subsequent cognitive development in infancy [9,10], thus interventions to promote motor skills have the potential to be used to enhance the overall infant development.

In the first part of this contingency study, the goal was for infants to discover and learn that the movements of a humanoid robot are contingent upon their movement. The robot performed a reward action (kicking a ball on a string) contingently, in response to a desired movement by the infant. Specifically, the robot rewarded the infant when s/he produced a leg movement above a specified, constant acceleration value, which we call the activation threshold. In the second part of this contingency study, we created a personalized contingency feedback adaption system that aims to encourage infants to gradually increase their peak acceleration of each movement.

This paper focuses on the evaluation of a reinforcement learning (RL) algorithm that moderates the adaptation of the activation threshold using the data distributions of the acceleration peaks of every infant from the first part of the contingency study. The experimentation presented here uses those data as input for the model, to generate activation threshold values that adjust to each distribution individually. This proof-of-concept of the model is a necessary step before carrying out a study with infants.

This paper is structured as follows: Sect. 2 presents related work from multiple fields. Next, Sect. 3 explains the origin of the infants' data from the first part of the contingency study, summarizing the foundational study that was carried out. Section 4 provides the details of the proposed model from the second part of the contingency study, from the discretization to build the set of thresholds to the RL-based approach. Section 5 presents a simulation of the model using the infant data. Finally, Sect. 6 summarizes the work and outlines next steps of this research.

2 Related Work

This multidisciplinary project brings together and builds on insights from multiple research areas. Section 2.1 describes the basic theory of infant motor development and the basis of contingency studies. Section 2.2 describes the importance of early intervention in atypical motor development and the need for personalized adaptation for each infant.

2.1 Infant Motor Learning and Adaptation

Current developmental theory proposes that infants learn the connection between their body and the environment by making frequent exploratory movements that help them to develop task-specific actions [1, 2]. For instance, when nine-month-old infants are placed in a jumper toy, they adjust the timing and force generation of their legs to optimize bouncing [11]. Our work used wearable inertial sensors attached to the infant’s limbs to track the acceleration and angular velocity of each limb throughout the motor exploration task.

To motivate infant movements, researchers use contingency feedback paradigms. Historically, infant contingency studies used a mobile paradigm where a specific arm or leg is attached to the mobile with a string. The more the infant moves the attached limb, the more sound and motion are generated by the overhead mobile [12]. Contingency studies have demonstrated that, when movements are reinforced by mobile motion, infants with typical development as young as three months old can increase the movement rate of the arm [13], increase the kicking rate of the leg [14, 15], move through a specific knee joint angle [16], produce more in-phase interlimb coordination by simultaneously moving both legs together [17], produce more in-phase hip-knee intralimb coordination by simultaneously extending the hip and knee of one leg [18], or produce selective hip-knee intralimb coordination (hip flexion with knee extension) by kicking a panel [19] or moving a foot vertically across a height threshold [20].

Those studies focused on reinforcing motion patterns; in this work we reinforce precise kinematic values, specifically the peak acceleration of a movement, aiming to encourage infants to increase the peak acceleration of their leg movements over time.

2.2 Infant Developmental Intervention

The main characteristic of this population is its enormous heterogeneity, since in such early stages, the aspects in the development and behavior patterns can vary enormously between individuals. That is why it is difficult to establish general guidelines and professionals need to make a more personalized analysis.

Approximately 9% of all infants in the United States are AR and could potentially benefit from early intervention services to address motor, cognitive, and/or social development [21]. All development domains, such as motor, cognitive, and social, are related, thus an intervention in one domain may provide benefits in all areas of development [15]. Despite this, the current standard of care for early intervention practice is to provide infrequent, low-intensity movement therapy or no intervention in infancy [22, 23]. New research has shown that early, intense, and targeted therapy intervention has the potential to improve neurodevelopmental structure and function [24]. Despite this potential gain, it can be challenging to find feasible and resource-efficient ways to deliver this type of intervention in infancy. Our proposed solution is to use a non-contact *socially assistive* humanoid robot to provide demonstrations and feedback aimed to encourage infants in movement exploration tasks. A key aspect of the efficacy

of this approach is the inclusion of *personalized models* appropriate for each infant participant that adapt the exploration task and difficulty to the specific infant, potentially allowing for higher engagement and improved learning.

Graded cueing is an approach that also aims at personalizing the level of task difficulty, by using increasingly specific cues or prompts given to the user [25]. This technique has been successful in rehabilitation of patients with brain injury and stroke, and has also been explored with socially assistive robots used with children with autism spectrum disorder in learning appropriate social skills [26,27]. The application of this technique consists of a set of steps that are applied sequentially. First, the therapist prompts the patient if the patient is having difficulties completing the assigned task. If, after a while, the patient continues to have difficulty, the therapist gives an increasingly specific cue, i.e., from a general verbal cue of patient’s body posture to a more specific cue such as imitating patient’s posture to help them to correct it. The purpose of using graded feedback is to encourage the patient to do most of the work on their own. The referenced past works address this problem by implementing models based on finite state machines or Markov Decision Processes. It has been shown to lead to more efficient learning and better learning outcomes.

This work follows a very similar concept. Different levels of difficulty are established and the participant starts at a low level. Difficulty levels are related to thresholds of acceleration peaks. The learning model must find the policy that allows to move between the different levels from the participant’s progress while maximizing the received reward (average acceleration). The idea is to adjust the specificity of the learning task – creating movements with higher acceleration – by adapting the acceleration threshold required to receive the contingency reward based on the infant’s past performance on the task.

3 Model Training Data

The training data used in this work were collected in a previous study. We summarize the data collection only briefly here.

Eight infants with TD between the ages of 6 and 8 months participated in a contingency feedback experiment in the Greater Los Angeles area. Only TD infants were recruited for this study as the first step was to enable the system to adapt to typical infant exploratory movement behavior.

The infant was placed in front of a NAO robot in a chair that allowed for full leg mobility, as shown in Fig. 1. The infant wore a head-mounted eye tracker. Opal inertial movement sensors [28] were affixed to each infant limb using cuffs with pockets. The sensors tracked the tri-axial acceleration and angular velocity of each limb.

For two minutes, the infant’s baseline movement was measured. During that time, the robot remained inactive. After the baseline, the robot demonstrated the reward action three times. The action was a basic knee flexion kick at a ball on a string. After the demo, the contingency phase of the study ran for eight minutes. If the infant produced an acceleration from the right leg above a




Fig. 1. An infant study participant interacting with the NAO robot in the previous study

fixed threshold of 3.0 ms^{-2} , the robot performed the reward action. We chose the acceleration threshold based on a previous study that measured the accelerations of infant leg movements [29]. In this study, the difficulty of the activity did not change and the threshold remained fixed throughout the session. The study was approved by the University of Southern California Institutional Review Board under protocol #HS-14-00911.

Table 1 shows the acceleration peaks from the eight infants in the study. The variance among the participants is notable. The values of the means vary based on performance during the session. For instance, infant 1’s mean peak acceleration is twice that of infant 5. Likewise, the maximum acceleration values reached by each infant and the number of acceleration peaks generated have a large variance. This is an indication that there is great heterogeneity in the participant pool, supporting personalized models rather than a generalized approach.

Table 1. Statistical outcomes of the study participants; N is the number of detected acceleration peaks for each participant.

VARIABLE	N	MEAN	STDEV	MIN	Q1	MEDIAN	Q3	MAX
ACC_PEAKS_U01	655	11.20	9.65	3.00	4.98	8.53	13.77	87.39
ACC_PEAKS_U02	417	9.77	8.06	3.01	4.30	6.31	12.51	45.66
ACC_PEAKS_U03	166	6.63	7.01	3.00	3.47	4.57	7.056	55.74
ACC_PEAKS_U04	326	9.51	8.61	3.02	4.21	5.87	11.15	63.49
ACC_PEAKS_U05	311	5.95	4.20	3.00	3.60	4.38	6.44	38.11
ACC_PEAKS_U06	499	8.98	8.69	3.00	4.20	5.78	9.38	72.41
ACC_PEAKS_U07	273	18.56	22.72	3.01	4.12	6.53	24.46	94.92
ACC_PEAKS_U08	359	7.11	6.16	3.01	3.85	4.98	7.88	48.26

The results of the previous study were promising and informed the objectives of this work. The majority of infants were able to learn the contingency with a set activation threshold. They moved above threshold more often in the contingency phase, in which they interacted with the robot, than in the baseline phase. Therefore, the next step is to try adjust the difficulty of the activity and determine if infants are able to adapt to a changing activation threshold.

4 User Adaptation Model

This section explains the proposed model for threshold adaptation in the infant movement contingency study. Section 4.1 provides a high level description of the problem. Section 4.2 explains the discretization of the peak acceleration values. Finally, Sect. 4.3 presents the RL approach for the adjustment of difficulty.

4.1 Problem Description

As noted earlier, the objective of the model is to adapt the activation threshold θ of the robot's reward action in real time. To achieve this, the contingency phase was segmented and the participants progress evaluated to determine the threshold for the next segment. Progress is defined in terms of the average of the acceleration peaks, since this work is focused on identifying thresholds that achieve a higher average in the acceleration of the infant's movements.

The threshold adaptation process was carried out during the contingency phase, in which the robot gave a reward (i.e., kicking the ball) each time the infant exceeded the current threshold, otherwise the robot remained still. Figure 2 is a representation of the contingency timeline divided into N segments. Each segment lasts 40 s; the duration was determined empirically to allow enough time for the infants to adapt to the new difficulty and for the model to receive enough learning experiences in every session.

The system started with an initial threshold θ^0 that changed over time based on the outcome obtained in each segment. At each time step n with $0 < n < N$, the model decides whether to raise, lower, or keep the threshold value θ^n , i.e., the difficulty of the activity (assuming higher thresholds are more difficult), based on the average value of the acceleration peaks obtained in the last segment. Each θ^n took its values from a set of thresholds Γ selected as described in Sect. 4.2.

The objective was to find the value of the threshold θ that maximized the acceleration of each infant's target limb. As shown in Sect. 3, the acceleration values reached by the infants are quite different from each other. Therefore, it is important to learn an individual model of each infant in order to obtain the threshold. The decision to modify the threshold is dependent on the threshold levels for each infant, the average acceleration value obtained in the previous segment, and the infant's degree of engagement. These variables were chosen because they are used by experts, and the aim is to learn a policy for each infant that adjusts the level of difficulty of the activity similar to the way a health care professional would.

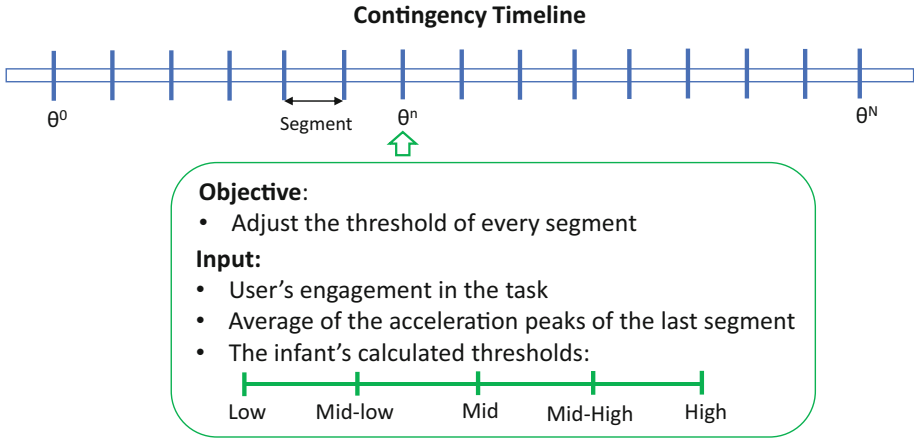


Fig. 2. Representation of the contingency problem

4.2 Discretization of the Acceleration Values

This section explains how the acceleration values of each infant were discretized to build a set $\Gamma = \{\theta_1, \theta_2, \dots, \theta_q\}$ composed of q discretized threshold values that best match the data collected in their past sessions. In this study, 5 levels of difficulty related to acceleration peaks were established a priori, i.e., $q = 5$. Additionally, we assumed Γ is sorted in ascending order, i.e., $\forall i, j$ and $i < j$, $\theta_i < \theta_j$ so that each threshold value corresponded to a level of difficulty: “low, mid–low, mid, mid–high, high”.

As discussed in Sect. 3, preliminary analysis of the data revealed large differences in the movement data captured from the participating infants; some demonstrated double the average acceleration peaks of others. This evidence is consistent with previous research in development [30]. Together with potentially higher variability within and across infants in different AR populations, this determined the need to create independent models for each participant. This, in turn, suggested that each infant should have a discretized set of thresholds, Γ , adapted to their abilities.

Instead of using a uniform discretization, we used a *K-means* algorithm with $k = 5$ that allowed for finding the five centroids that best separated the acceleration data for each infant [31]. The centroids were directly related to the five levels of difficulty of the problem. Therefore, each threshold value $\theta_i \in \Gamma$ corresponded to a different centroid. Figure 3 shows an example for the data gathered from infant 1. The graph is the representation of the allocation of the instances to the different clusters found by the algorithm (the blue points corresponds to the instances in cluster 1, the green points to the instances in cluster 2, and so on). Furthermore, each cluster is represented by a centroid that corresponds to a value associated with the level of difficulty (in this case, $\Gamma = \{4.97, 10.81, 17.32, 28.89, 52.56\}$). In this example, and in most of the participants, there is no homogeneous allocation of the instances in the clusters due

to the way in which the data are distributed: 47% (low), 29% (mid-low), 15% (mid), 6% (mid-high), 2% (high) for the infant 1. This means that most instances are concentrated around low levels of acceleration, since infants reach the highest peaks of acceleration at specific times.




Fig. 3. Estimation of thresholds of the infant 1 using K-Means for the discretization of the accelerations peaks

4.3 Mapping the Threshold Adaptation Problem onto Reinforcement Learning

In this section, we describe the mapping of the problem of threshold adaptation of an infant described in Sect. 4.1 onto an RL approach. Such modeling requires defining all the elements of a Markov Decision Process (MDP): the state and action spaces and the reward and the transition functions [32]. We consider this to be an episodic task, where for each episode the infant is evaluated in N steps.

In this work, a state $s \in S$ is a tuple in the form $s^n = \langle \xi^n, \theta^n \rangle$, where ξ^n and θ^n are respectively the disengagement of the infant and the current threshold of the system at step n . Feature ξ is a binary feature, i.e., $\xi \in \{0, 1\}$, where $\xi = 0$ if the infant is engaged, and $\xi = 1$ otherwise. Instead, feature θ takes values from the discrete set $\Gamma = \{\theta_1, \theta_2, \dots, \theta_q\}$ built by discretizing the acceleration values of each infant, as described in Sect. 4.2. Therefore, the size of the state space S is $2 \times q$.

In state s^n , the agent performs an action $a^n \in A$. We consider the action space A as being composed of three actions, $A = \{-1, 0, 1\}$. These actions are used to decrease, leave as is, or increase, respectively, the threshold θ^n of the current state.

After performing an action a^n in state s^n , the agent transits to a new state $s^{n+1} = \langle \xi^{n+1}, \theta^{n+1} \rangle$. A transition function is required to compute the values for ξ^{n+1} and θ^{n+1} . The value of ξ^{n+1} is computed using Eq. 1:

$$\xi^{n+1} = \begin{cases} 1, & \text{if } countHits < 2. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

where *countHits* is the number of times the infant moves with an acceleration above or below threshold θ^n in step n . To compute the value of θ^{n+1} , we assume that $\theta^n = \theta_i$, i.e., θ^n at step n corresponds with the i -th threshold in Γ . Then, we compute θ^{n+1} as in Eq. 2.

$$\theta^{n+1} = \theta_{i+a^n} \quad (2)$$

Therefore, if $a^n = 1$, the threshold is increased and θ^{n+1} takes the value of the $(i + 1)$ -th element in the Γ set, i.e., $\theta^{n+1} = \theta_{i+1}$. Conversely, if $a^n = -1$, the threshold is decremented and takes the value of the $(i - 1)$ -th element, i.e., $\theta^{n+1} = \theta_{i-1}$. If it is unchanged, then $\theta^{n+1} = \theta_i$.

Finally, when the learning agent performs an action a^n in a state s^n and moves to a state s^{n+1} , it also receives a reward signal r^n . We formulate the reward function as shown in Eq. 3.

$$r^n = \begin{cases} 0, & \text{if } countHits = 0. \\ avgSuccAcc \times (countSuccHits/countHits), & \text{otherwise.} \end{cases} \quad (3)$$

where *avgSuccAcc* is the average acceleration of the infant’s movements above threshold θ^n , *countSuccHits* is the number of times the infant moves with an acceleration above the threshold θ^n , and *countHits* is the number of times the infant moves (above or under the threshold θ^n). The rationale behind the reward function in Eq. 3 is as follows. If the infant does not move, the reward received is 0. If the infant moves ($countHits > 0$), and the threshold θ^n is exceeded ($countSuccessHits > 0$), the reward is greater than 0. If the threshold is easily exceeded by the infant, the reward is expected to be higher, consistent with a higher threshold. Conversely, if the threshold is not easily exceeded by the infant, the reward decreases, since *countSuccessHits* tends to 0.

Finally, the reward function in Eq. 3 is different from the reward the robot provides to the infant. The former is used to learn a policy by RL to regulate the threshold θ that best fits the infant, while the latter is used to motivate the infant every time the infant exceeds the current threshold.

5 Simulation Evaluation of the Model

This section describes the evaluation of the model from Sect. 4. The objective is to extensively test the model prior to a study with infants by using the data from the first part of the contingency study described in Sect. 3. In a real scenario, the discretization of the acceleration values would be done individually

from the data of the past sessions of each of the infants. To train the model, the peaks of infant movement acceleration were simulated and used as input to the RL model. Acceleration peaks of the participants typically followed an exponential distribution where there was a higher concentration of instances at low accelerations and fewer at high accelerations, as can be shown in Fig. 4. From the calculated distributions of each of the infants, the system generates random acceleration values that follow this distributions. In this way, it can be said that the behavior of every infant was being imitated, in terms of acceleration, based on their past experiences.

The objective of this simulation evaluation was to test the behavior of the model with two completely different infants: infant 5 and infant 7. According to Fig. 1, infant 5 obtained an average peak acceleration of 5.956 with a maximum value of 38.101, while infant 7 obtained an average peak acceleration of 18.56 with a maximum value of 94.92. Although they were very different, both followed an exponential distribution, as can be seen in Fig. 4. After applying the discretization described in Sect. 4.2, the set of thresholds for infant 5 were $\Gamma = \{3.41, 4.97, 7.88, 12.21, 20.87\}$ while the those for infant 7 were $\Gamma = \{3.82, 7.58, 16.55, 31.77, 56.99\}$. Both sets presented different values in line with the outcomes of each infant.




Fig. 4. Graphs of the distributions of the acceleration peaks of infants 5 and 7.

The simulation followed the approach presented in Fig. 2, in which the phase of contingency was divided into steps. Every step was an experience for the model, in which the values of acceleration peaks were created from the distribution of each infant, see Fig. 4.

For each infant, we simulated 50 episodes of 20 steps as described in Sect. 4.3. We used Q-Learning, and ϵ -greedy as exploration-exploitation strategy [32]. Table 2 shows the resulted Q-tables for infants 05 and 07 at the end of the learning process. Between state S_0 and state S_4 are the states when the infant was engaged with the task, i.e., $\xi = 0$, while from state S_5 to state S_9 , the infant was disengaged, i.e., $\xi = 1$.

Significant differences can be seen between the Q-tables. The values of infant 7 are higher than those of infant 5, since the episodes generated with the first one

contributed with greater rewards than those of the second. Looking at the highest value of each row, the policy learned can be found for each participant. For infant 5, the resulting policy considers more adequate to stay in a low threshold. For instance, if we consider that the infant starts in state S_0 , the best action is to stay in this state (the action *stay* is the action with the highest Q-value). However, if the infant starts in state S_3 and is never disengaged, the policy decides to transit first to S_2 , then to S_1 and, finally, to S_0 , the state with the threshold that best suits the infant. As a last example, if we consider the infant starting in state S_4 , the system could transit to state S_8 (i.e., system transits from a state with a high threshold to a state with a mid-high threshold, although the infant has disengaged). In S_8 , the best action is to reduce the threshold, so that the system could move to S_2 (if we assume the infant is engaged again), then to S_1 , and finally to S_0 . Finally, it is important to note that the rows for the states S_5 and S_6 are all 0. This is intuitive since the infant is never disengaged when the system is in low or mid-low threshold values and, hence, the states S_5 and S_6 would be never visited.

Table 2. Results of Q tables of the simulated experiments of infants 05 and 07.

	USER 05			USER 07				
	UP	STAY	DOWN	UP	STAY	DOWN	(deng/th)	
S0	86.34	88.98	0	337.42	335.97	0	(0/Low)	Engaged
S1	84.09	87.76	89.30	350.06	348.44	345.34	(0/Mid-Low)	
S2	80.27	85.77	87.05	330.02	335.30	341.83	(0/Mid)	
S3	73.27	75.74	86.03	300.29	322.43	342.31	(0/Mid-High)	
S4	0	0	32.22	0	301.57	318.51	(0/High)	
S5	0	0	0	0	0	0	(1/Low)	Disengaged
S6	0	0	0	0	0	0	(1/Mid-Low)	
S7	0	61.03	0	0	0	0	(1/Mid)	
S8	68.47	78.44	83.34	0	0	0	(1/Mid-High)	
S9	0	70.93	78.50	0	275.99	313.25	(1/High)	

Instead, infant 7 was able to get higher acceleration values between mid-low to mid thresholds, since s/he had higher accelerated movements in past sessions. Following the same reasoning as in the other Q-table, if we consider the infant that starts in state S_0 and is never disengaged, the system first transits to S_1 , and then to S_2 . Then, a loop occurs: in S_3 it decides to reduce the threshold and transits to S_2 . Therefore, the threshold that best suits this infant is between the states S_2 and S_1 . Finally, as in the previous case, the rows for states S_5 to S_8 are 0, as these states are never visited; this infant does not disengage until reaching a high threshold value in state S_9 .

6 Conclusion

This paper presented an approach for using a personalized reinforcement learning algorithm for infants learning to reach target leg movement acceleration. The RL-based model was able to determine the best threshold configuration in terms of peak acceleration. The results of the simulation were very promising; the model was sensitive to the high variance among the infant study participants. The policy learned for each participant indicated the thresholds that would reach higher rewards values. Since the reward function was related to the average of the acceleration peaks and the number of peaks detected, maintaining these thresholds in a session would help to maximize these two variables.

In a real infant-robot interaction scenario, higher difficulty levels would offer better rewards from the robot. Thus, the ultimate goal of this study is to determine whether the robot is able to encourage the infant to reach higher accelerations from their movements to get better rewards from the robot. This work validates the proof-of-concept of the model, making it ready for implementation in our upcoming contingency study of infant-robot interaction.

This novel work in socially assistive robotics for infant movement therapy is the basis for the upcoming studies that will extend the presented results. We plan to explore new reward functions that reinforce other aspects of the movement or allow the dissociation of one limb from the other. Additionally, we intend to integrate this socially assistive robot system into the next infant-robot contingency study to determine if the model helps with the adaptation of the infants achieving better results than with approaches based on fixed activation thresholds.

References

1. Gibson, E.J., Pick, A.D.: *An Ecological Approach to Perceptual Learning and Development*. Oxford University Press, Oxford (2000)
2. Thelen, E., Smith, L.: *A Dynamic Systems Approach to the Development of Cognition and Action*. The MIT Press, Cambridge (1994)
3. Smith, B., Vanderbilt, D.L., Applequist, B., Kyvelidou, A.: Sample entropy identifies differences in spontaneous leg movement behavior between infants with typical development and infants at risk of developmental delay **5**, 55 (2017)
4. Kouwaki, M., Yokochi, M., Kamiya, T., Yokochi, K.: Spontaneous movements in the supine position of preterm infants with intellectual disability. *Brain Dev.* **36**(7), 572–577 (2014)
5. Rademacher, N., Black, D.P., Ulrich, B.D.: Early spontaneous leg movements in infants born with and without myelomeningocele. *Pediatric Phys. Ther.* **20**(2), 137–145 (2008)
6. Smith, B.A., Teulier, C., Sansom, J., Stergiou, N., Ulrich, B.D.: Approximate entropy values demonstrate impaired neuromotor control of spontaneous leg activity in infants with myelomeningocele. *Pediatr. Phys. Ther.* **23**(3), 241–247 (2008)
7. McKay, S.M., Angulo-Barroso, R.M.: Longitudinal assessment of leg motor activity and sleep patterns in infants with and without down syndrome. *Infant Behav. Dev.* **29**(2), 153–168 (2006)

8. Geerdink, J.J., Hopkins, B., Beek, W.J., Heriza, C.B.: The organization of leg movements in preterm and full-term infants after term age. *Dev. Psychobiol.* **29**(4), 335–351 (1996)
9. Kermoian, R., Campos, J.: Locomotor experience: a facilitator of spatial cognitive development. *Child Dev.* **59**, 908–917 (1998)
10. Oudgenoeg-Paz, O., Volman, M.: Attainment of sitting and walking predicts development of productive vocabulary between ages 16 and 28 months. *Infant Behav. Dev.* **35**, 733–736 (1998)
11. Goldfield, E.C., Kay, B.A., Warren, W.H.: Infant bouncing: the assembly and tuning of action systems. *Child Dev.* **64**(4), 1128–1142 (1993)
12. Rovee-Collier, C.K., Gekoski, M.J.: The economics of infancy: a review of conjugate reinforcement. In: *Advances in Child Development and Behavior*, vol. 13, pp. 195–255. Elsevier (1979)
13. Watanabe, H., Taga, G.: General to specific development of movement patterns and memory for contingency between actions and events in young infants. *Infant Behav. Dev.* **29**(3), 402–422 (2006)
14. Heathcock, J.C., Bhat, A.N., Lobo, M.A., Galloway, J.: The performance of infants born preterm and full-term in the mobile paradigm: learning and memory. *Phys. Ther.* **84**(9), 808–821 (2004)
15. Lobo, M.A., Galloway, J.C.: Assessment and stability of early learning abilities in preterm and full-term infants across the first two years of life. *Res. Dev. Disabil.* **34**(5), 1721–1730 (2013)
16. Angulo-Kinzler, R.M., Ulrich, B., Thelen, E.: Three-month-old infants can select specific leg motor solutions. *Motor Control* **6**(1), 52–68 (2002)
17. Thelen, E.: Three-month-old infants can learn task-specific patterns of interlimb coordination. *Psychol. Sci.* **5**(5), 280–285 (1994)
18. Angulo-Kinzler, R.M.: Exploration and selection of intralimb coordination patterns in 3-month-old infants. *J. Motor Behav.* **33**(4), 363–376 (2001)
19. Chen, Y.-P., Fetters, L., Holt, K.G., Saltzman, E.: Making the mobile move: constraining task and environment. *Infant Behav. Dev.* **25**(2), 195–220 (2002)
20. Sargent, B., Schweighofer, N., Kubo, M., Fetters, L.: Infant exploratory learning: influence on leg joint coordination. *PLoS ONE* **9**(3), e91500 (2014)
21. Rosenberg, S.A., Robinson, C.C., Shaw, E.F., Ellison, M.C.: Part c early intervention for infants and toddlers: percentage eligible versus served. *Pediatrics* **131**(1), 38–46 (2013)
22. Roberts, G., Howard, K., Spittle, A.J., Brown, N.C., Anderson, P.J., Doyle, L.W.: Rates of early intervention services in very preterm children with developmental disabilities at age 2 years. *J. Paediatr. Child Health* **44**(5), 276–280 (2008)
23. Tang, B.G., Feldman, H.M., Huffman, L.C., Kagawa, K.J., Gould, J.B.: Missed opportunities in the referral of high-risk infants to early intervention. In: *Pediatrics peds-2011* (2012)
24. Holt, R.L., Mikati, M.A.: Care for child development: basic science rationale and effects of interventions. *Pediatr. Neurol.* **44**(4), 239–253 (2011)
25. Bottari, C., Dassa, C., Rainville, C., Dutil, E.: The IADL profile: development, content validity, intra- and interrater agreement. *Can. J. Occup. Ther.* **77**(2), 345–356 (2009)
26. Feil-Seifer, D., Matarić, M.: A simon-says robot providing autonomous imitation feedback using graded cueing. In: *Poster paper in International Meeting for Autism Research (IMFAR)* (2012)

27. Greczek, J., Kaszubski, E., Atrash, A., Matarić, M.: Graded cueing feedback in robot-mediated imitation practice for children with autism spectrum disorders. In: IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 561–566 (2014)
28. APDM Wearable Technologies, Portland, OR, USA, Opals. <https://www.apdm.com/wearable-sensors/>. Accessed 15 July 2018
29. Trujillo-Priego, I.A., Smith, B.A.: Kinematic characteristics of infant leg movements produced across a full day. *J. Rehabil. Assist. Technol. Eng.* **4**, 2055668317717461 (2017)
30. Adolph, K.E., Robinson, S.R.: Sampling development. *J. Cogn. Dev.* **12**(4), 411–423 (2011). <https://doi.org/10.1080/15248372.2011.608190>
31. Hartigan, J.A., Wong, M.A.: Algorithm as 136: a k-means clustering algorithm. *J. Roy. Stat. Soc. Ser. C (Appl. Stat.)* **28**(1), 100–108 (1979)
32. Sutton, R.S., Barto, A.G.: Reinforcement Learning I: Introduction. MIT Press, Cambridge (1998)