Principles for Building "Smart" Learning Environments in Pediatric Early Rehabilitation

Elena Kokkoni*, Ashkan Zehfroosh*, Prasanna Kannappan*, Effrosyni Mavroudi†, James C. Galloway‡, Jeffrey Heinz§, René Vidal†, and Herbert G. Tanner*

*Mechanical Engineering, University of Delaware, Newark, DE 19716
†Biomedical Engineering, Johns Hopkins University, Baltimore, MD 21205
‡Physical Therapy, University of Delaware, Newark, DE 19713
§Linguistics and Cognitive Science, University of Delaware, Newark, DE 19716

***Abstract*—Young children with motor delays lack early mobility rehabilitation applicable to natural and complex environments. The goal of this work is to create a novel "smart" learning environment that combines socially assistive robots and rehabilitation technology to promote mobility early on. This paper outlines the main principles the "smart" environment is built upon, the main challenges for creating it, and presents preliminary data on the feasibility of every component.**

I. Introduction

Motor impairment in the early stages of life significantly also affects cognitive, and social development [1]. Yet, for young children at risk, there is a dearth of appropriate, personalized pediatric mobility rehabilitation paradigms capable of increasing the dosage of mobility and applicable to natural and complex environments. The overarching goal of the research effort summarized here is the development of such a paradigm. Motivated by evidence of benefits seen in other areas of motor rehabilitation with older children [8, 4], the envisioned paradigm blends active and adaptive socially assistive robotic technologies with innovative and age-specific passive elements, such as a portable body-weight support device, to boost mobility dosage within enriched environments.

There are great challenges associated with the development and implementation of such paradigms. One of them is the lack of adequate evidence in tests with very young populations, that would support the hypothesis that social interaction between children and machines can be expected. Other challenges relate to the development of the rehabilitation system itself. For example, no complete theory for adaptive social child-robot interaction has been established to guide the design of the system and the adjustment of its parameters.

Driven by the need to address some of these challenges, this work brings together kinesiology, robotics, computer and cognitive science to create a novel "smart" early rehabilitation environment. This is the first study to combine socially assistive robots and other rehabilitative devices into a single paradigm. The aspects of this work are reported here. First is the development and integration of several technological components into a proof-of-concept single robot-assisted early rehabilitation paradigm. The second is the assessment of the capacity of this environment for early pediatric motor rehabilitation, and its potential for broader future use.

II. Challenges

The paper identifies three major challenges:

- What is the most appropriate learning environment to advance rehabilitation outcomes?
- Is Child-Robot Interaction (CRI) viable at a very young age? Can we use CRI to motivate very young children to explore learning environments?
- Can we build "smart" robots in order to have an integral role in complex learning environments used in early pediatric rehabilitation?

A. Learning environments supporting rehabilitation outcomes

The environment plays an important role into shaping behavior: enriched environments that contain high levels of complexity and novelty have shown to induce experience-based brain and behavioral changes when explored in high-dosages [6]. The challenge in this case is that very young children with motor disabilities often do not have the capacity to explore their environment due to both limited motivation and mobility impairment. The amount of experiences they can gain throughout their lifetime is therefore severely more limited compared to typically developing children.

B. Viability of CRI

The fundamental hypothesis that drives this work is that an early start in rehabilitation is critical. In other words, to achieve better results in the rehabilitation outcomes, one starts as early as possible in life, when most brain changes typically occur and the onset of most developmental milestones is observed [3]. CRI has not been adequately explored with children younger than two years of age, and thus there is no evidence of guaranteed social interaction. In addition, there is no standard theory to guide the design and configuration of CRI systems for very young children. For instance, it is not clear what mathematical models are appropriate, or what CRI decision-making architectures are effective.

C. Robots in complex pediatric rehabilitation environments

Complex learning environments promise greater benefits. Complex environments, however, are "noisy;" they are considered unpredictable and allow less control authority. That is where one of the major challenges to robot automation comes about. We need to build robots that are adaptive; robots that automatically and safely interact with young children in complex environments, not only in the structured environment of a research lab but also in the clinic, home, or school. Automation involves the robots being able to recognize children's behavior and respond. This process has to be accurate, fast, and safe, otherwise children lose interest and interaction diminishes. To have real impact, such automation technology has to be portable, low-cost, and capable of handling the "noise" in natural home settings—because this is precisely where high-dosage of training can be achieved.

III. Approach

A number of steps are taken to approach the three aforementioned challenges. A *dynamic* rehabilitation environment is designed, where both the child and the robot interact and learn from each other.

A. The learning environment

The environment is enriched with large objects to explore while working on balance, coordination, and strength. Because the manipulation of such objects require actions that are beyond the children's level of ability, a body-weight support device (OASUS™; Enliten LLC) is added to ease mobility and exploration (Fig. 1). Body-weight support devices are used in traditional rehabilitation paradigms as a way to alter the dynamics of gait and behavior [5, 8].

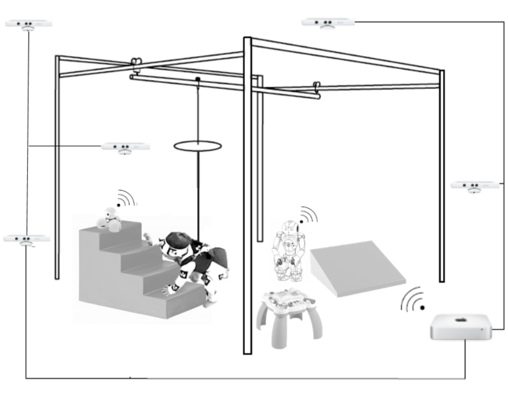

Fig. 1. Learning environment and equipment setup for early pediatric motor rehabilitation.

B. Early CRI setup

Two robots, a humanoid (NAO™; Aldebaran Robotics) and a wheeled toy robot (Dash™; Wonder Workshop), are chosen as dynamically adaptive and real-time controllable elements of the system. The role of the robots is to motivate the child to manipulate objects and perform complex tasks. This is done through either social interaction or actual manipulation (e.g. NAO mimicking human motor actions such as hand manipulation of table-top surfaces and walking to a goal; Dash moving fast horizontally and/or climbing inclined surfaces).

Proximity is a way to view the success of such interactions. For example, if the distance between the robot and the child is decreased or remains the same while, say, Dash is climbing the incline, this may be an indication that the child is actually following the robot and performs the task as well. (EK: connect proximity with the MDP later!)

C. Robot behavior in complex learning environments

Our approach focuses on realizing two main robot functionalities.

1) Feedback—receiving information about the environment and child's actions: A data acquisition and analysis system is designed to feed (potentially real-time) information about the child's actions and objects in an environment. Both types of information are critical for the robot's decision-making: how to respond to changes in the environment, how to direct social interaction, and how to manipulate objects.

Different types of data are acquired and integrated using the Robot Operating System (ROS). These data involve synchronized video from a network of five Kinect sensors, and tracking coordinates of AR tags placed on the child's body. All the above are used for the child's activity recognition.

Activity recognition in a particular session is supported by two separate analyses based on the data collected in previous sessions. The first is a manual, off-line annotation process of the child's activities from pre-recorded video. The second, is the training and tuning of machine learning activity recognition algorithms based on the annotated data. The goal is for the configured algorithms to be able to classify activities in the streaming video of a session in (almost) real-time.

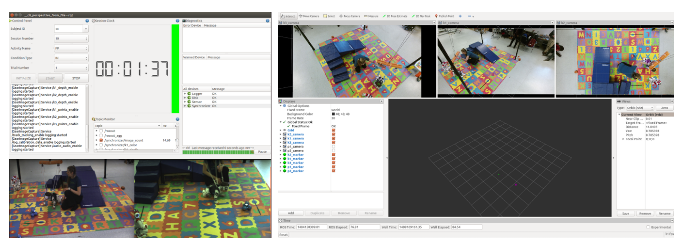

Fig. 2. Screenshots of the graphical user interface of the software environment used for data acquisition and preliminary analysis.

2) Decision-making—response of the robots to developments in the scene: Activity recognition enables robot situational awareness which in turn supports decision-making. The robots decide on what is the most appropriate action on their part to promote the rehabilitation outcomes, e.g. keep the child in motion for as long as possible. The decisions are based on the history of past CRI experiences, and are the product of yet another machine learning and (discrete-event) optimal control algorithm. This algorithm is based on a crude mathematical model that attempts to capture human behavior through a Markov Decision Process (MDP). In the MDP model,

BT: Recent ACC paper models humans using MDPs states are the possible configuration of the child and the robot. Action set of the MDP model is defined by different activities of the robot and the child that cause transitions among states. The fact that direct control is only over the robot's actions motivates us to consider activities of the child as outcomes of such actions in a probabilistic manner. Therefore, we drop the child's actions from the action set of the MDP and make the transitions probabilistic instead. The system is initiated with some prior regarding those transition probabilities and a machine learning algorithm updates the transition probabilities in real-time through observations during the course of CRI via maximum likelihood notion. The machine learning algorithm that updates those model parameters needs to produce effective estimates fast and with very small data. For this reason, the learning module utilizes a technique used primarily in natural language processing called smoothing [2] over the emprical observations.

IV. Validation

A two-year old child with Down syndrome participated in eight sessions that took place twice a week for four weeks. The child was able to crawl but not walk. Each session lasted for one hour. Information on child's mobility and interaction was used to validate our approach at all three levels.

A. Feasibility

Outcome: The system was used successfully by the child and the environment afforded him significant opportunities for exploration.

Specifically, the child explored the objects of the learning environment and used the body-weight support device at all eight sessions (Fig. 3). Indeed, the child was able to perform all complex tasks, like climbing an inclined platform and ascending the staircase. The assistance of the body-weight support device promoted the onset of new motor behaviors that had not been observed before.

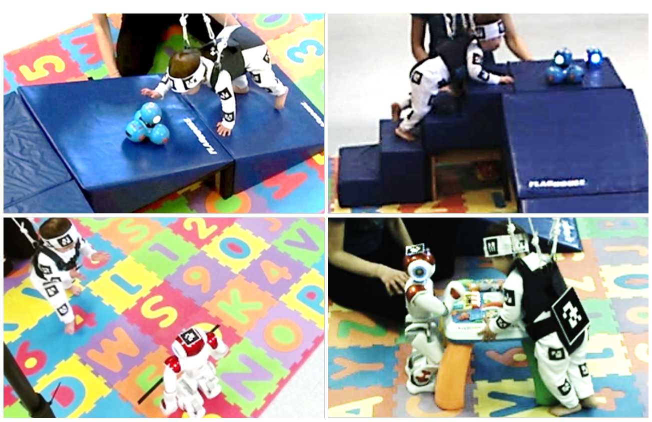

Fig. 3. Infant is performing complex motor tasks.

B. Effectiveness

Outcome: The child's exploration was due in a great part to his interaction with the robots.

During all sessions, the child interacted with both robots most of the time, but in different ways. Specifically, the subject was engaged into a game of chase aimed to increase his mobility, but he seemed to do so more with Dash than with NAO (Fig. 4). NAO afforded manipulative actions by the child that were cause-and-effect (e.g., kicking and pushing to throw the robot down and see it standing back up) and more complex in nature (e.g., kicking requires coordination of more body segments and maintaining balance).

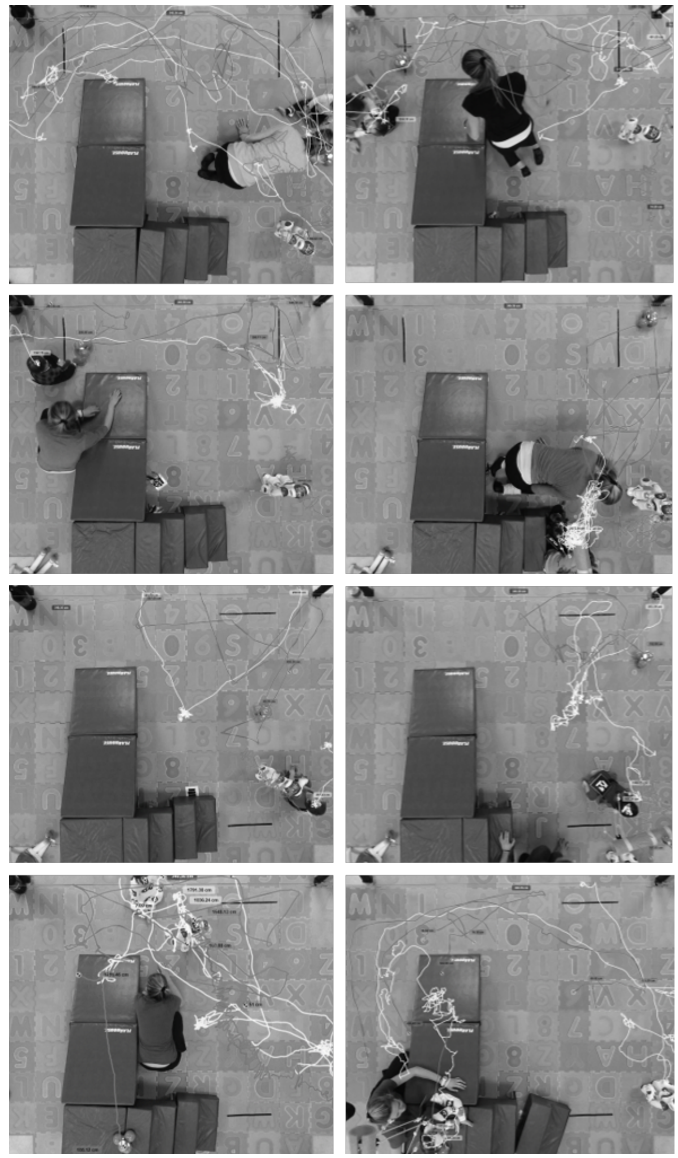

Fig. 4. The child's movement path (shown in white) in the environment while interacting with the robots. Most of his mobility can be attributed to the chasing of Dash. Left column pictures involve trials of the child moving without the assistance from the body-weight support device whereas right column with the assistance of the device. Each row of pictures is from a different session in the series of sessions.

C. Potential

Outcome: Adaptive robots may be potentially used in complex learning environments.

1) Feedback: Environment and child behavior information was successfully collected from all Kinect sensors in real-time. The child's actions are currently used to train machine learning algorithms that will be used to identify overall system state and support robot control. However, consistently receiving tracking coordinates from the AR tags has been challenging for two main reasons: first, the size of the tags on the subject's limbs was too small to provide adequately regular measurements, and the child frequently removed the straps where the tags were attached to. The latter challenge was later addressed through a special "tag suit," worn by the subject on the last session (see Fig. 3). The former challenge persists. Even with the tag suit, tracking the tags on the limbs has been problematic, whereas the bigger tags placed on the subject's trunk were considerably easier to track. Alternative solutions using motion capture systems have been explored without significant success primarily due to occlusions. Work on the limb tracking issue is ongoing.

As far as child action recognition is concerned, we present preliminary results on action classification from a sample of annotated videos ($N = 13$ videos) using established algorithms [9]. In particular, we extract appearance and motion features along dense trajectories, as shown in Fig. 5, we then encode feature descriptors by aggregating first and second order statistics using a Fisher Vector [7] and finally classify the videos using linear Support Vector Machine (SVM) classifiers. The actions of interest are associated with the major motor milestones that typically developing children acquire during their first year of life: sitting, crawling, standing, and walking. These are actions that the system needs to recognize—ideally in real-time. For example, if the child spends too much time in a stationary position such as sitting, the system should recognize that, and decide on an action that will elicit a mobility response from the child, such as crawling—since the target is to keep children moving and exploring.

We trim the videos into segments, with each segment containing one of the four actions of interest. The final number of instances per action were: 6 for crawling, 15 for sitting, 18 for standing, and 23 for walking. With available video footage from five cameras, the total number of video segments available for training and testing the action classification method is 310. Note that since these videos were collected in complex natural environments, occlusions, multiple entities (people and robots) simultaneously moving, and viewpoint variations make the task of action classification challenging. To this end, we split the videos into five different training and testing groups in two different ways. First, we randomly use 80% of the videos belonging to each class as training data, and leave the rest for testing. In the second way, we randomly split videos into training and testing sets as before, but now making sure that all videos corresponding to a single action are in the same group (either used for training or testing). The average action classification accuracy over 5 random splits generated using the first method is 72.6%, while the second method yielded 63.1%. The action classification accuracy is expected to be better in the first setup, since it is easier for the algorithm to classify a testing video in case it has seen the same behavior instance from another view during training. Still, in both cases the action classification accuracy is above chance level.

Figure 6 illustrates the confusion matrices for two random splits. In this instance, the figure shows the action classifier being unable to recognize the crawling, since the number of annotated samples was very small (just 30 videos corresponding to six different crawling instances). Future work will examine the success of the same algorithms to accurately recognize actions from bigger samples.

Fig. 5. Visualization of dense trajectories [9] extracted from a video sample. Red dots correspond to trajectory positions in the current frame. Each moving point is tracked for 15 frames.

2) Robot decision-making: The learning component of the CRI framework that was set up has shown promise in terms of adaptation and customization to different human subjects, in support of an future automation loop closure. We used an MDP model for human behavior as a response to robot action, in the context of a basic "game of chase" between robot and child. An instantiation of such an MDP model is shown in Fig. 7. Robot actions are modeled as transitions. The transition label set is $\{f, s, b\}$, with f representing a decision that the robot makes to move "forward" and toward the child, s is a decision to stay still, and b expresses a decision to move backward. States are labeled with two letters; the first describes the process that the robot is in after making a decision labeled in $\{f, s, b\}$, and the second represents how the child is responding to the ensuing robot behavior. Symbol F stands for the robot moving forward (toward the child), S denotes the state where the robot stands at a constant distance from the child, and B represents the condition where the robot moves backward and away from the child. Labels G and N mark the child's response, with G expressing the fact that the child is making progress toward the rehabilitation goal—in this case, to remain in motion—and N marks the absence of progress. Positive rewards will thus be assigned to the states that have a G component in their label, and negative or zero rewards will be assigned to the remaining ones.

The MDP model is originally initialized with a transition graph in which the robot actions can only trigger deterministic transitions between states that preserve the second state label.

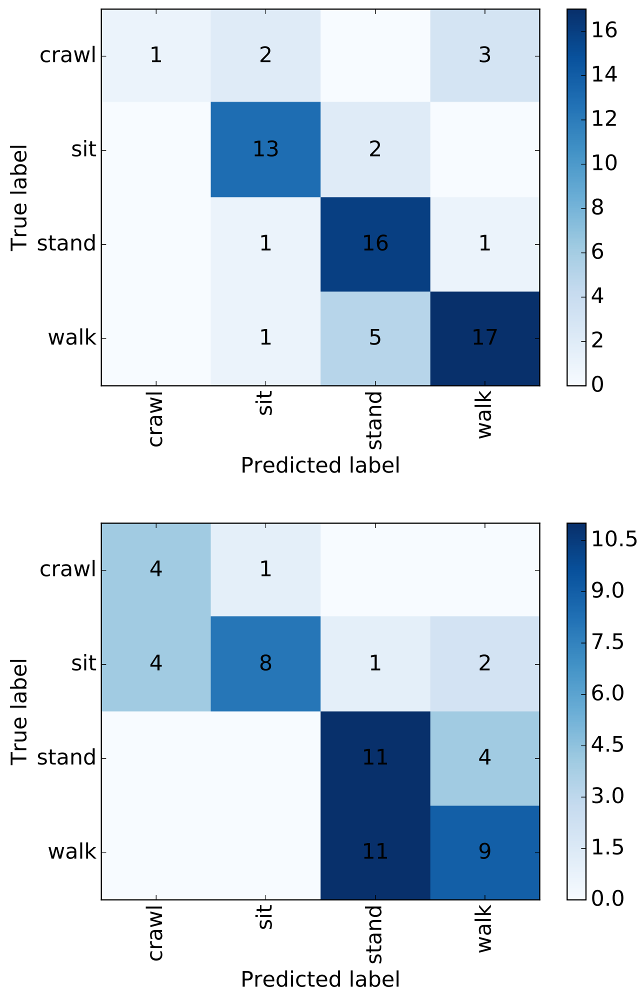

Fig. 6. Action classification preliminary results. Diagonal entries of confusion matrices show the number of correctly classified action instances per action class with respect to ground truth annotations. (Upper): Results for a random training/testing split generated by using 80% of video samples for training and the rest for testing. (Lower): Results for random training/testing split using 80% of video samples for training and the rest for testing, with videos corresponding to different views of the same action instance being used in the same set.

For example, action b at the state labeled (F, N) is causing a transition to state (B, N). Thus the MDP of Fig. 7 is initialized with the blue transitions only. But when a robot action is *observed* to trigger a transition, say from state (F, N) to state (F, G)—the red arrow in Fig. 7—then transition f at state (F, N) becomes nondeterministic, and different transition probabilities are assigned to blue and red jumps according to empirical data. This process will in principle continue until the evolving MDP converges to some true hypothesized probabilistic model of this particular CRI.

A serious challenge in this application space, however, is that the set of available observations is typically very small, and as a result, a naive method for updating transition probabilities, say Maximum Likelihood Estimation (MLE), will produce a poor MDP model. For this reason, smoothing algorithms [2] are adapted and applied to this problem, as an attempt to overcome the scarcity of learning data. Simulation results with the aforementioned model have so far supported the hypothesis that smoothing can be advantageous as a learning mechanism for populating the unknown parameters of this MDP.

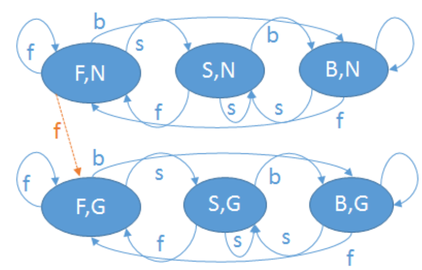

Fig. 7. Example of our MDP model.

V. Conclusions

This paper describes the design and validation of a novel robot-assisted learning environment that has the potential to be used in future pediatric rehabilitation. As part of the development of such an environment, a remotely actuated playground has been developed to allow complex and rich interaction between children, toys, and robots. In parallel, a visual sensor network has been implemented and operated with a dual purpose: to record data for assessment of rehabilitation outcomes, and for providing (hopefully, real-time) feedback for the automation system to guide the robot in support of these rehabilitation objectives. for the purpose of The initial realization of this combined sensing and computation framework has shown promise with (a) preliminary evidence of affording exploration by the child, and (b) data suggesting the viability of the framework for early CRI.

The ultimate goal is to enable high-dosage pediatric rehabilitation in natural and complex environments that could take place outside the structured setup of an academic lab or clinic. We envision "smart" environments that are robot-assisted but not human-sterile. The intention is not for the automated CRI system to become the sole form of interaction for the young child. Instead, it is envisioned that judicious rehabilitation environment designs can serve as catalysts for peer-to-peer and other forms of (human) social interaction.

Acknowledgments

This work has been supported by NIH #5R01HD87133.

References

[1] Joseph J. Campos, David I. Anderson, Marianne a. Barbu-Roth, Edward M. Hubbard, Matthew J. Hertenstein, and David Witherington. Travel Broadens the Mind. *Infancy*, 1(2):149–219, 4 2000. ISSN 15250008. doi: 10.1207/S15327078IN0102{_}1. URL http://doi.wiley.com/10.1207/S15327078IN0102_1.

[2] Stanley F. Chen and Joshua T. Goodman. An Empirical Study of Smoothing Techniques for Language Modeling.

(June), 1996. ISSN 08852308. doi: 10.3115/981863.981904. URL http://arxiv.org/abs/cmp-lg/9606011.

[3] Sharon E Fox, Pat Levitt, and Charles A Nelson. How the timing and quality of early experiences influence the development of brain architecture. *Child development*, 81(1):28–40, 2010. ISSN 1467-8624. doi: 10.1111/j.1467-8624.2009.01380.x.

[4] Maninderjit Kaur, Timothy Gifford, Kerry L Marsh, and Anjana Bhat. Effect of Robot Child Interactions on Bilateral Coordination Skills of Typically Developing Children and a Child With Autism Spectrum Disorder : A Preliminary Study. pages 31–37, 2013.

[5] Max J Kurz, Wayne Stuberg, Stacey Dejong, and David J Arpin. Overground body-weight-supported gait training for children and youth with neuromuscular impairments. *Physical & occupational therapy in pediatrics*, 33(3):353–65, 8 2013. ISSN 1541-3144. doi: 10.3109/01942638.2013.771719. URL http://www.ncbi.nlm.nih.gov/pubmed/23477615.

[6] Jess Nithianantharajah and Anthony J Hannan. Enriched environments, experience-dependent plasticity and disorders of the nervous system. *Nature reviews. Neuroscience*, 7(9):697–709, 9 2006. ISSN 1471-003X. doi: 10.1038/nrn1970. URL http://www.ncbi.nlm.nih.gov/pubmed/16924259.

[7] Florent Perronnin, S Jorge, and Thomas Mensink. Improving the Fisher Kernel for Large-Scale Image Classification To cite this version : Improving the Fisher Kernel for Large-Scale Image Classification. 2010.

[8] Laura a Prosser, Laurie B Ohlrich, Lindsey a Curatalo, Katharine E Alter, and Diane L Damiano. Feasibility and preliminary effectiveness of a novel mobility training intervention in infants and toddlers with cerebral palsy. *Developmental neurorehabilitation*, 15(4):259–66, 2012. ISSN 1751-8431. doi: 10.3109/17518423.2012.687782. URL http://www.ncbi.nlm.nih.gov/pubmed/22670679.

[9] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. *Proceedings of the IEEE International Conference on Computer Vision*, pages 3551–3558, 2013. ISSN 1550-5499. doi: 10.1109/ICCV.2013.441.