A Simple Model of Skill Acquisition in a Dynamic Balance Task

Ruta Desai* and Jessica K. Hodgins*
* Carnegie Mellon University, USA
{*rutad, jkh*}*@cs.cmu.edu*

1 Motivation

We are interested in modeling human motor skill acquisition in complex dynamic tasks. Understanding how humans learn to do a task can unveil new directions for learning of control in robotics. Modeling motions of human learners can also help us in generating effective motions for graphics and entertainment robotics. Human motor adaptation has been widely studied in structured tasks such as goal directed arm movements, to answer questions such as 'what is learned' and 'how is it learned' [1]. However, it is unclear whether the models proposed using studies of such structured tasks apply to highly dynamic motor skills such as balancing and locomotion. We study motor skill acquisition in a motion capture friendly, highly dynamic task of balancing on a bongo-board.

2 Approach

Four subjects (Three Females, One Male) with no prior experience of balancing on a bongo-board, tried to learn the task of balancing on a bongo-board in a motion capture laboratory for three days alternatively with rest days (Fig. 1A). During this time their actions were captured using 16 infrared cameras of Vicon motion capture system (120 Hz.). An average of 10 takes of 2 minutes each, were captured during an hour long capture session each day. Subjects were allowed to use the board only during the experiment time. All four subjects showed tremendous improvement in balancing on the bongo-board (evaluated using average balance time on the board (Fig. 1C)).

2.1 PCA based Learning Discovery

We hypothesize that humans use basic primitive motions that combine together and enable the performance of complex tasks such as balancing on bongo-board. During the learning process, subjects explore different weighted combinations of these primitive motions. Weights of the primitive motions that are less useful for the task reduce as the subjects learn. Decomposing the subject's motion into such basic primitive motions and evaluating contribution of individual primitives (using weights) to the overall motion can thus provide a model of what the subjects are trying to learn. However, the set of primitive motions necessary for balancing on bongo-board are not known. We assume that the motions of the subjects on the final day (last day of experiments) hold the information about the necessary primitive motions for the task of bongo-board balance. We use Principal Component Analysis (PCA) [2] to decompose subjects' motion on the final day into basic primitive motions. For a given set of data, PCA finds orthogonal linearly independent basis vectors or bases (the eigen vectors) such that the vector space spanned by them captures the maximum variance of the data. The PCA bases obtained from the motion data thus correspond to most significant motions which capture the variation in the given data. We use these motions corresponding to the PCA bases as the primitive motions for our task.

Figure 2 (left) shows these primitive motions (P_i) obtained from the final day data of all subjects using PCA. Centered lower body joint angle data of all subjects from all takes of final day was used for this purpose. We only use lower body data (pelvis and below) as subjects seem to sparingly use upper body after the first few takes. We use the first 5 PCA bases as primitive motions because they account for 95% of the variance of the data. The net motion M can then be described as a weighted combination of these primitives (Eq. 1).

$$M = \sum_{i=1}^{5} w_i P_i \tag{1}$$

where w_i denotes the weight corresponding to the primitive P_i. We assume that the primitive motions obtained from the final day data also account for the variance of the data from day-1 and day-2 for all subjects. This is a simplifying assumption. In reality, the subjects might be simultaneously learning the primitives P_i and their weights w_i. If so, it would be hard to compare a subject's motion across different days of learning and hypothesize about the changes in control as (s)he learns. Using our assumption, we fix the primitives P_i and calculate their corresponding weights for day-1, day-2 and day-3 motions, for individual subjects. We average these weights over all takes for a particular day for individual subjects. These mean weights of each primitive motion are compared over the period of learning for each subject (Fig. 2 right). The resultant trends of how these weights change as subjects learn give an insight into what the subjects are learning (See Sec. 3).

2.2 Linear Control using Identified Primitive Motions

We are also interested in creating control simulations which mimic the motions of subjects learning to balance on a bongo-board. Instead of tracking reference trajectories [3] obtained from motion capture to create such motions in simulation, we use a linear control approach called Output Feedback Control (OFC) inspired by recent success of such approaches in balancing a simple model over variety of surfaces [4]. This is because we are interested in obtaining a control which inherently generates different motions of the human learners instead of a control that tracks a fixed target motion. This approach also allows us to answer whether we can modify well-established controllers to produce human-like motions. OFC provides an easy way of integrating the primitive motions identified by PCA analysis in the form of outputs used for feedback. We model the lower body in 2D using a four-bar linkage with telescopic links (Fig. 1D). The wheel is modeled as a cylinder. All contacts are modeled as rolling contacts with no slip. The model has five degrees of freedom (DOF) in total. The ankle, hip and knee joints (modeled as telescopic joints) are actuated. We linearize this model to obtain a linear system representation (Eq. 2).

$$\dot{x} = Ax + Bu, \qquad y = Cx \qquad (2)$$

where x is the state vector, u is the control input vector and y is the output vector. OFC finds a linear control law $u = -Fy$ to stabilize this system. Similar to Nagarajan and Yamane [4], the time-invariant output feedback gain matrix F is found using the convergent iterative algorithm in [5]. This algorithm requires an output matrix C, a state gain matrix Q and a control input gain matrix R. While Nagarajan and Yamane [4] select outputs intuitively, we use primitives obtained using PCA as outputs. The output feedback matrix C is thus composed of PCA bases or eigen vectors. State gain matrix Q is constructed by diagonalizing the corresponding eigen values of the PCA bases. Constructing the Q matrix in this way allows us to emphasize primitives preferred by the human subjects while obtaining F. Intuitively, this means that we measure our outputs and take feedback in the space defined by the primitive motions used by the human subjects.

3 Results

Subjects explore various primitive motions (P_1-P_5) such as rocking in the lateral and sagittal planes and yaw about the pelvis as they learn to balance on the bongo-board (Fig. 2 right). The decrease in weights w_5 of pelvis yaw P_5 suggests that subjects learn to control their yaw in order to improve balancing on the board (Fig. 2E). A weak trend of decrease is also seen in the two coupled primitive motions of rocking P_3 and hip flexion-extension P_4 in sagittal plane suggesting improved control in sagittal plane (Fig. 2C,D). Two distinct control strategies for balancing on a bongo-board (rocking vs. non-rocking) emerge out of the trends of weight changes w_1 of lateral rocking primitive P_1 (Fig. 2A). Subject-1 and Subject-3 decrease rocking on day-2, followed by an increase on day-3 (see w_1). While rocking on day-1 may correspond to falling behavior and instability on the board, the decrease on day-2 suggests improved control. The increased w_1 on day-3 suggests use of controlled rocking as a strategy to balance on the board. Subject-2 and Subject-4 instead show an increase in rocking on day-2 and further a decrease on day-3. Thus they seem to explore rocking motion on day-2 and then choose against using it on day-3. The correlation between weights w_1 and w_2 also suggest the use of knee and ankle flexing to increase the range of motion of the board for generating lateral rocking (Fig. 2A,B).

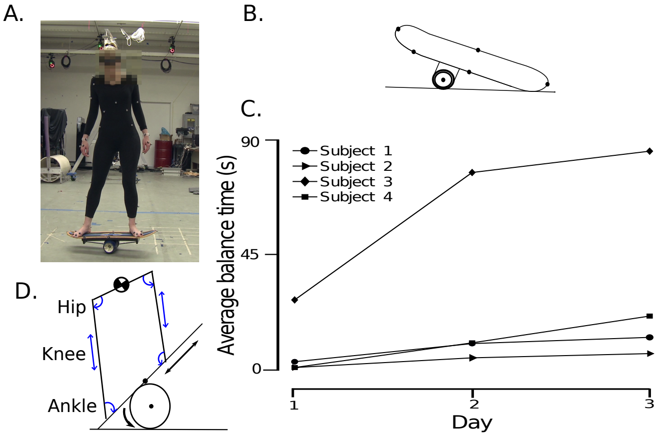

Figure 1: Experiment setup and performance evaluation: (A) A subject wearing standard set of 60-markers, balancing on a bongo- board during motion capture. (B) Bongo-board motion is also measured with 7 markers (black dots). (C) Subjects improve their balancing ability which can be measured by average balance time on the board over the days of experiment. (D) A simplified 2D model of human balancing on a bongo-board in simulation. The model has 5 DOF (unactuated board has 2 DOF: rotation of the wheel and sliding of the board (black arrows)). All hip, knee and ankle joints are actuated (blue arrows).

As a first step towards a control simulation which produces motions like those of human learners, we used OFC with outputs chosen as primitive motions and state gains chosen as eigen values of those motions. We find that without any manual tuning, these outputs and gains allowed us to get a stabilizing controller which could balance the non-linear 2D model (Fig. 1D) in an ideal simulation (without any noise or modeling errors). However, the resultant motions do not visually look human like. Although we do not aim for robust controllers, the range of attraction (ROA) for the resulting controllers measured as the maximum board angle deviation that can be recovered from, is also much smaller than what we see in humans (max. ROA: 0.5° (control), 15° (humans)).

4 Future Work

We hypothesized that humans use weighted combination of basic primitive motions for complex motor tasks and used PCA to see how such primitives evolve as subjects learn a dynamic task of balancing on a bongo-board. The emergence of distinct strategies for control (rocking vs. non-rocking) from our analysis suggests the need of more data (more subjects, motion capture our longer learning interval beyond 3 days). Our preliminary simulation using linear Output Feedback Control modified to use human data, could stabilize a simplified 2D model of human balancing on a bongo-board. However, the resultant motion did not compare to human motion. A more detailed simulation (with 3D model, accounting for human like noise and delay) might be necessary to simulate motions of human learners.

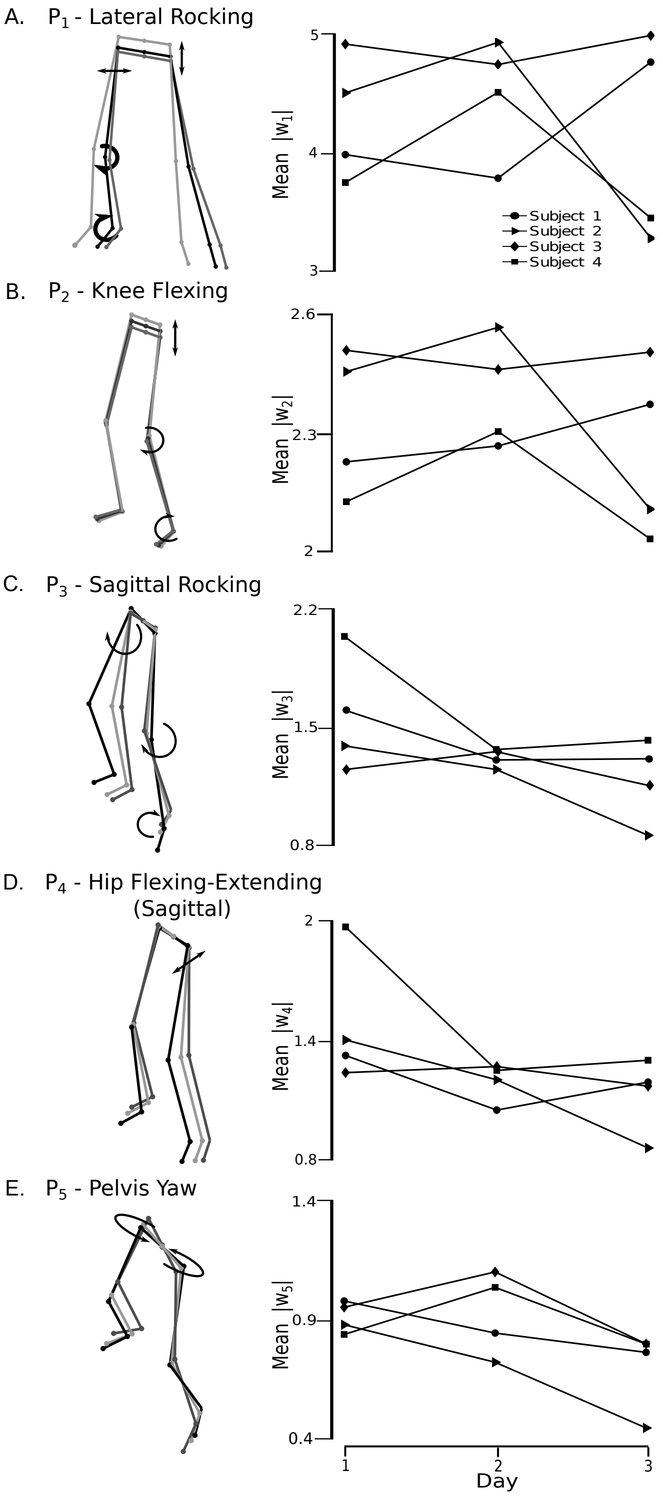

Figure 2: Primitive motions and their corresponding weights. (A, B) Primitive P_1 is rocking in lateral plane using hip, knee and ankle joints while P_2 is bending using knee and ankle joints. (C, D) P_3 and P_4 primitives correspond to sagittal plane motions: rocking in sagittal plane using hip, knee and ankle joints and flexion-extension motion of hip joints respectively. (E) P_5 is yaw about the pelvis. The right column shows how the weights for each of the five primitives vary over the three days among the different subjects.

References

[1] D. M. Wolpert, J. Diedrichsen and J. Randall Flanagan, "Principles of sensorimotor learning", Nature Reviews Neuroscience, 2011.

[2] `http://en.wikipedia.org/wiki/Principal_component_analysis`

[3] K. Yamane and J. K. Hodgins. "Simultaneous tracking and balancing of humanoid robots for imitating human motion capture data", IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems, 2009.

[4] U. Nagarajan and K. Yamane, "Universal Balancing Controller for Robust Lateral Stabilization of Bipedal Robots in Dynamic, Unstable Environments", IEEE Intl. Conf. on Robotics and Automation, 2014.

[5] J. Yu, "A convergent algorithm for computing stabilizing static output feedback gains", IEEE Trans. on Automatic Control, 2004.