2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

Extensive Human Training for Robot Skill Synthesis: Validation on a Robotic Hand

Erhan Oztop[12], Li-Heng Lin[2], Mitsuo Kawato[12], Gordon Cheng[12]
[1]*JST, ICORP, Computational Brain Project, 4-1-8 Honcho Kawaguchi, Saitama, Japan*
[2]*ATR Computational Neuroscience Laboratories, 2-2-2 Hikaridai, Seika-cho, Soraku-gun Kyoto, 619-0288, Japan*
{erhan, mlin, kawato, gordon }@atr.pj

***Abstract* – We propose a framework for skill synthesis for robots that exploits the human capacity to learn novel control tasks. The conceptual idea is to incorporate the target robotic platform into the experimenter's body schema so that it can be controlled effortlessly as if the robot were a part of the body. Once this stage is achieved, the dexterity on a task exhibited with the new external limb –the robot- can be used for designing controllers for the task under consideration. This article exemplifies the proposed framework by showing the derivation of an effective open-loop controller that can manipulate two balls with the fingers of a 16-DOF robotic hand.**

Index Terms –Body Schema, Skill Synthesis, Hand Control

I. Introduction

As soon as we grab a computer mouse, it becomes a part of our body; we control it effortlessly and fluently as we control our hands. This introspection has a neural basis[1]. Recent neurophysiological experiments with behaving monkeys has shown that primates are endowed with a very plastic representation of their arms and hands, which are expanded instantaneously as soon as a tool is grabbed that can be utilized to manipulate the space[2, 3]. The cortical representation of one's body is generally referred to as *body schema*. Accumulating evidence suggests that the body schema is very plastic and parts of it can be swapped in or out in a modular fashion depending on the context and task requirements.

We propose that the adaptability of human body schema provides a smart and effective way of programming humanoid robots that are required to be equipped with robust and dexterous skills. The scenario envisioned involves a humanoid robot that is completely operated (through a sophisticated intuitive interface) by an expert human controller, where the learning software of the robot builds controllers by monitoring sensor readings and motor commands generated by the human controller. The proposal has parallels with learning by demonstration and imitation learning [4-11]; however, the bottle neck of these systems, i.e. the large difference between the structure of the demonstrator and the learner is overcome by human *motor learning*. In particular, we do not promise an easy task for the human teacher, on the contrary, we anticipate long learning periods (e.g. weeks, months) after which the robot learning can take place.

In this article, we report our first step towards validation of this proposal using a 16-DOF robotic hand, where a two-ball manipulation task was selected as the target skill. We are contend that the success of this study will lend support for our proposal as the task is complex and the robot and human hand have large differences in terms of kinematics, dynamics and elasticity. If we can achieve our goal then the second stage will be to develop algorithms for deriving autonomous controllers for the tasks controlled by human (motor) intelligence. In an earlier report, we have presented the preliminary results of this stage, where the robot could swap balls very slowly (7.5 seconds/swap)[12].

In this article, we focus on how this basic skill can be improved in terms of smoothness, speed and to what extent the improvement can be made. In addition, for completeness, we review the experimental setup and the task in the following sections.

II. The Task and The Hardware

A. The Ball Swapping Task

The target task for our proposal of using human visuomotor learning for robot skill synthesis was chosen as the manipulation of so called Chinese healing/health balls. The task is defined as the manipulation of the balls such that the initial positions of the balls are swapped (see Figure 1).

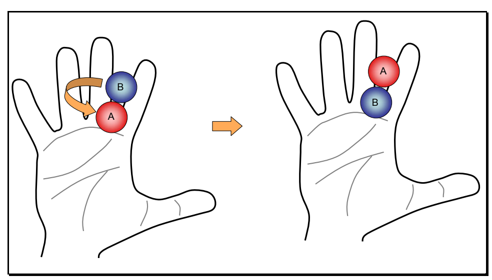

Figure 1. The ball swapping task. The goal is to swap the position of the balls without dropping them.

Humans can perform this task in several ways at speeds of 1-2.5 Hz depending on experience, often requiring palm and thumb articulation. Awkward it may be, the rotation can also be carried out with no palm and thumb articulation at much lower rates (< 1 Hz). From this outset it was not possible to predict whether the task can be completed with the Gifu Hand, our robotic platform that we describe next.

B. The Target Robotic Platform

The Gifu Hand III (Dainichi Co. Ltd., Japan) consists of a 4 degrees of freedom (DOF) thumb and four 3-DOF fingers.

The Gifu Hand is connected to a PC hosting A/D, D/A, Counter and Timer PCI boards. The A/D PCI cards allow the PC to read the motor currents. The Counter PCI cards are used to obtain number of encoder clicks, i.e. joint angle changes. The D/A cards convert the PC's digital outputs to analog voltages that are amplified and sent to finger motors. The PC runs a PD controller at 500Hz.

For the ball swapping task, the Gifu Hand was mounted on PA-10 robot arm (Mitsubishi Heavy Industries), which enabled us to adjust the orientation of the Gifu Hand. The hand was oriented so to have gravity apply enough force on the balls producing a rolling motion towards the finger tips. This adjustment was performed once and kept fixed through out the study. Finally, we constrained ourselves to four finger manipulation since, (because of technical reasons) the thumb was not available during the initial experiments.

C. Human Motion Capture

Human motion capture is performed using the Visualeyez (PhoeniX Technologies Inc.) that is an active marker based tracking system. Visualeyez comes with a soft glove that allows user adjustable positioning of the markers (see Figure 2). Visualeyez system has an accuracy around 1mm in 3D position; however the system intrinsically suffers from occlusions. In addition, fast motion of the fingers degrades the accuracy of the motion capture.

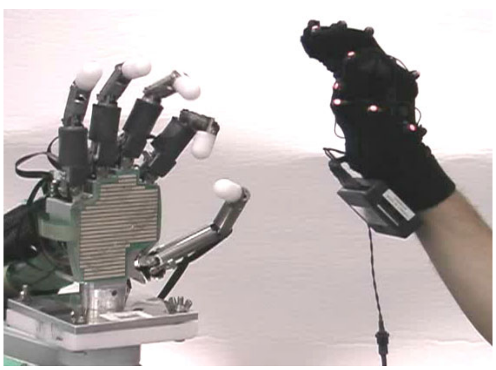

Figure 2. The real-time human control via visual data capture. Subject finger motions are tracked at 30 Hertz, converted into joint angles and sent to the PD controller as desired joint angles.

D. System Implementation

The integrated control system consists of multiple components which run independently and communicate via UDP (see Figure 3). The first component encapsulates the Visualeyez system, which captures and sends out user finger tip positions in real-time (at 30Hz). The tip positions arrive at Inverse Kinematics block in which the desired endeffector positions are converted into desired joint angles (see section Human Robot Control for details) that are sent to the Central Controller. The Central Controller receives all system information and controls all system components. A User Interface is hooked to the Central Controller which provides the user the facilities to see the received system information (e.g. current joint angles and motor currents) and control the system manually (if desired). Finally, the Gifu Hand Controller receives desired joint angles from the Central Controller and runs a PD controller at 500Hz to move the fingers to the desired joint angles.

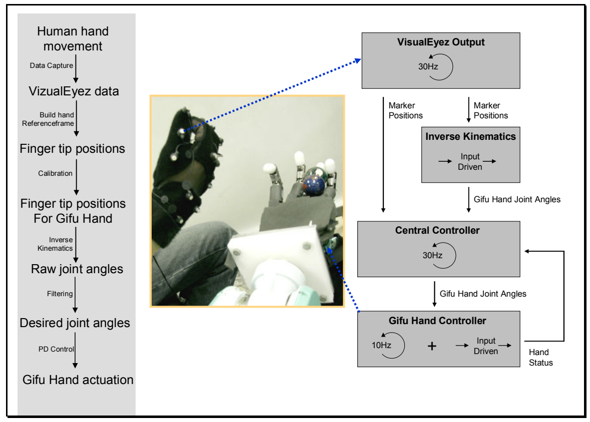

Figure 3. Left: Logical flow of the robot control via the human operator. Right: Actual implementation of the Visualeyez–Gifu Hand real-time control system.

III. Human Robot Control

To minimize the difficulty of learning the ball swapping task for humans an intuitive control of the Gifu Hand is necessary. The straightforward approach is to attach enough number of sensors to finger tips and knuckles so that the finger joint angles can be calculated using the geometry formed by the sensor positions in 3D space. This approach, however, produced poor results in terms of intuitive control of the robot hand. We have noted two reasons for this: Firstly, the markers placed on knuckles were not stable due the stretch of the glove when fingers moved. Secondly, the differences in the structure and capability of human and robot hands were not accounted with this approach. Humans actuate their palms in order to help fingers achieve their goals, as in pinching. In contrast, the Gifu Hand has no palm actuation capability. So a pinching movement at the human side would not produce pinching at the robot side, which renders the control unintuitive. Therefore we used finger tip positions as targets for the Gifu Hand fingers, requiring calibration and inverse kinematics steps (described in the next subsection.) This way we avoided the problem with knuckle motion, and compensated for the human palm actuation by robot finger motion. Three sensors were attached at the base of the hand to construct a reference frame for representing the finger tip positions. So a total of seven sensors were used (three for base and four for fingers, see Figure 4)

A. Mapping Human Finger Positions to the Robot Hand

The reference frame constructed on the hand (*hand reference frame* here forward) is chosen to roughly match the reference frame defined on the robot. However, it is not possible to guarantee that the coordinate frames will be aligned exactly, as each time the Visualeyez glove is worn the marker locations would slightly vary. For this purpose, a calibration procedure was introduced to align the human hand finger tip positions with the robot hand finger tip positions. The postures shown in Figure 4 are used to collect 16 data points for data fitting. The goal of calibration is to find a 4x4 calibration matrix ($\mathbf{M}$) to represent the points given in hand reference frame ($\mathbf{H}$) in the robot reference frame ($\mathbf{G}$) so that $\mathbf{HM} = \mathbf{G}$ holds. Denoting the points in hand reference frame by bar notation, **H** and **G** is given by

$$\mathbf{H} = \begin{bmatrix} \bar{x}_1 & \bar{y}_1 & \bar{z}_1 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ \bar{x}_n & \bar{y}_n & \bar{z}_n & 1 \end{bmatrix}, \mathbf{G} = \begin{bmatrix} x_1 & y_1 & z_1 & 1 \\ \vdots & \vdots & \vdots & \vdots \\ x_n & y_n & z_n & 1 \end{bmatrix} \text{(with n=16).}$$

The (approximate) solution is obtained with pseudo inverse as $\mathbf{M} = \mathbf{H}^{\dagger}\mathbf{G}$ so that given, say, human index finger tip position $\begin{bmatrix} \bar{x} & \bar{y} & \bar{z} & 1 \end{bmatrix}$ the desired index tip position in the robot reference frame would be given by $\begin{bmatrix} \bar{x} & \bar{y} & \bar{z} & 1 \end{bmatrix}\mathbf{M}$.

Figure 4. The calibration postures used are shown. The finger tip sensor positions marked with dashed circles are used as data points for calibration (only the first posture is marked).

B. Finger Inverse Kinematics

Each finger used in this study has 3 DOF forming a non-redundant effector chain except at the singularities. Since the finger DOFs are relatively low, for inverse kinematics we have simply constructed a lookup table that can be efficiently indexed by the desired end-effector position [13]. The table entries were constructed by slicing the lateral abduction-adduction joint into 70, and the remaining two flexion joints into 90 bins. This required approximately 13Mbytes of memory (with float = 4 bytes) which is not a problem for normal desktop computers. Each finger of the Gifu Hand has the same kinematics defined in finger specific reference frame, so we did not need to duplicate the table for each finger. The lookup table we adopted was a KD-tree which allowed very efficient searches, which proved to be faster than iterative inverse kinematics solution.

IV. Human Training and Robot Skill Acquisition

A set of guidelines for human learning emerged while carrying out the study, which were effective for the ball swapping task. We present these guidelines without quantitative analysis in the current report.

A. Human Training

Observation of the motion of the robot fingers while the subject moves his/her own fingers for several hours, suffices to achieve an intuitive human control of the robot fingers; however, to control the robot for a manipulation task requires longer training. The subject has to discover the affordances [14] available to the robotic hand, learn to combine them gracefully building a set of motor primitives [8, 10], and finally use these primitives in junction with the visual feedback (i.e. the position of the balls) to achieve the task goal. Note that subject has to learn these without any somatosensory feedback which is crucial for everyday manipulation.

Our preliminary experiments showed that the following stages were effective for visuomotor learning of the ball swapping task, which conform the theory of motor skill learning [15].

1. Move fingers without any ball and observe the robot
2. Manipulate single ball without dropping it
3. Fine tune (2) to have control over the motion of the ball by rolling it back and forth between index-middle finger and little-ring finger apertures.
4. Manipulate two balls without dropping
5. Devise a strategy for completing the task
6. Fine tune (5) for completing the task

Steps 1 and 2 is usually achieved in hour, the skill of Stage 3 is best attained after one day of motor consolidation [16-19]. Stage 4 and 5 develops rather in parallel; different strategies are explored while playing with the balls. Stage 6 requires visual monitoring of the balls and the robot fingers in coordination, hence takes longer to achieve.

B. Autonomous Ball Swapping on the Robot

Imitation and learning from observation/demonstration research in robotics [4-11, 20, 21] aims at removing the burden of robot programming from the experts by letting non-experts 'teach' the robots. The simplest method to transfer a skill to a robot is to directly copy the motor commands of the demonstrator to the robot, so called the 'motor tape' approach [22], which is extremely effective for open-loop tasks. However, in general it is not possible to adopt this approach. First, motor commands may not be available to the robot; even if available, the differences between the demonstrator and the robot often render the motor commands useless for the robot.

The situation in our case, however, is different; because the 'correct motor commands' on the robot is produced by the human operator/experimenter. For this convenience, the price one has to pay is the effortful human training that may not only include controlling the robot but also accomplishing a 'hard' task as well. In other words, instead of expert robot programming we rely on subjects' motor learning ability to produce the motor commands on the robot itself, which can be conveniently played back. If the task is open-loop controllable then the motor tape approach produces an easy and robust way for skill synthesis. Having said this, it should also be noted that designing an intuitive interface can be time consuming. On the other hand, it can be argued that this design phase can be standardized and so, need not be repeated for each new skill to be synthesized.

V. Results: Subject and the Task Execution

A. Subjective Assessment of Training Experience

For this study, one subject (author E.O.) went through the training procedure. It took two days to learn the stages 1-5, by spending several hours per day. However, the lack of tactile feedback rendered Stage 6 difficult, making it longer for the subject to master the final stage. Even though attempting to swap balls without dropping was initially frustrating, the task turned out to be learnable. In fact, it became 'easy' for the subject after a week of training and the control did not require subject's conscious effort. After the training, the subject gained full control of the robot hand and could easily detect the affordances provided by the hand at each session. Each time the data capture glove is worn, the calibration procedure must be reapplied. During this procedure the subject may produce the calibration postures in slightly different ways. This results in slight differences in the control afforded by the particular calibration, which could be detected by the subject.

The key strategy (motor primitive) observed for ball swapping was 'to kick' the ball rather than sliding it over the surface of the fingers (see Figure 5). This is very different from what humans do when the task is executed in natural conditions. In the natural setting tactile feedback guides the execution where the visual component of the task is minimal. It is notable that the kicking strategy emerged without conscious awareness during the Stages 2-3.

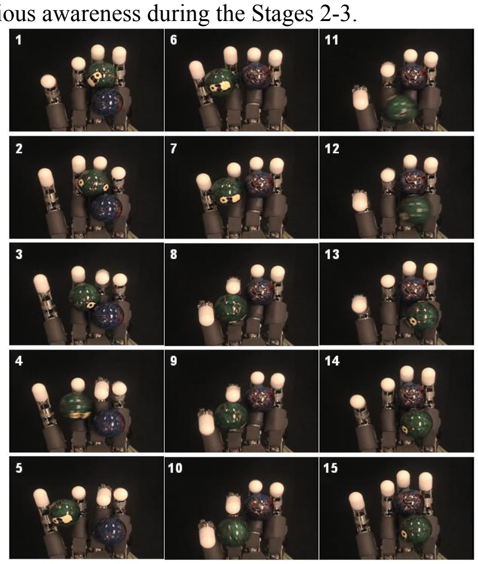

Figure 5. Frames representing the ball swapping task performed by the Gifu Hand using the skill transferred from the human performance. A movie clip showing the execution is available at http://www.cns.atr.jp/~erhan/Movies/ball1-short.mpg.

B. Ball Swapping Task Performance

We recorded several successful ball-swapping performed by the subject at 30 hertz. Without any filtering we have played back one of the performances on the Gifu Hand. The only data manipulation was to crop the portion of the trajectory at the end where no movement was executed, and to linearly interpolate it to the beginning so that we could playback the performance in a looped fashion. The performance was good, albeit slow, as the subject had executed the task in a slow and controlled manner (see Figure 5). The robot could swap the balls without dropping for long periods (>30 minutes). The sensitivity to initial conditions of the balls was minimal. Many types of disturbances were tolerated and self-corrected with this open-loop controller obtained through human dexterity. Moreover, although we have not used different balls during human performance, the controller was able to swap balls with different sizes and weights, such as a wooden and coated metal ball; a coated but larger metal ball, and finally a large but very light Styrofoam ball [12] .

Inspection of the joint angle trajectories revealed that subject used several kicks to roll the ball over the fingers. One of these was apparent when the actual performance of the subject or the playback on the robot was observed, where the little and ring fingers were raised and a kick was applied to roll the ball towards the index finger (the ellipse in Figure 6 indicates this motion). The other rather smaller kicks were noticeable neither to the observers, nor to the subject himself prior to inspection of the trajectories.

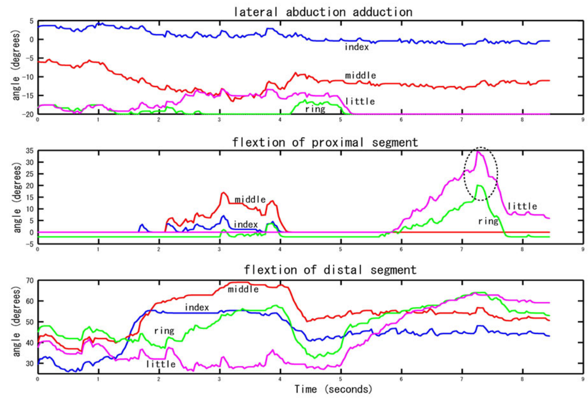

Figure 6. Joint angles produced on the Gifu Hand by a successful ball swapping performed by the subject. Horizontal axis denotes time in seconds; vertical axis denotes angles in degrees. Top panel: The lateral abduction-adduction angles. Middle panel: distal joint angles (the Gifu Hand's distal two joints are coupled and have the same angle). Bottom panel: proximal joint angles (the angles at the base of the fingers). The curves are label with 'little', 'ring', 'middle' and 'index' indicating the associated fingers.

V. Results: Improving Robot Performance

At this stage of the project, we focused on improving robot performance that was obtained via human visuomotor learning. The next stages will address the issues how the robot can self-improve, or learn in an interleaved manner with the human.

The natural candidates for improvement was the elimination of noise and improvement of the execution speed, as the robot trajectory produced by the human on the robot was slow and noisy due to device and human factors.

A. Smoothing

For smoothing, a Gaussian kernel fit was used. N Gaussians ($\sigma^2 = 0.1$) with equally spread means over the time axis are used to generate an N dimensional vector on which a linear regression is applied. Let **T** be the duration of a single ball swapping motion then let

$$\mathbf{k}(t) = \begin{bmatrix}\phi_1(t) & \phi_2(t) & \cdots & \phi_N(t)\end{bmatrix} \text{ where } \phi_j(t) = e^{-\frac{(t-t_j)^2}{\sigma^2}},\ t_j = \frac{j}{N}T.$$

Denoting the Gifu Hand joint angle at time t with the row vector **θ(t),** we can perform the linear fit as follows

$$\mathbf{W} = \mathbf{K}^{\dagger}\mathbf{A} \text{ where } \mathbf{K} = \begin{bmatrix}\mathbf{k}(0)\\ \mathbf{k}(t_i)\\ \vdots \\ \mathbf{k}(T)\end{bmatrix} \text{ and } \mathbf{A} = \begin{bmatrix}\boldsymbol{\theta}(0)\\ \boldsymbol{\theta}(t_i)\\ \vdots \\ \boldsymbol{\theta}(T)\end{bmatrix}$$

Then the joint angles at time t can be approximated by the weighted sum of Gaussian kernels as $\tilde{\boldsymbol{\theta}}(t) = \mathbf{k}(t)\mathbf{W}$. The value of N determines the smoothing applied (note that σ^2 was kept constant). A large N results in over fitting with noise retained. Experimentation with different parameters indicated that with smooth trajectories the robot failed to perform the ball swapping, apparently because the impact of the kicks were washed out by the smoothing.

B. Smoothing with Impulse Preservation

As noted earlier, the critical primitive used by the subject for ball swapping was kicking. At certain critical points in the trajectory the kicks (impulses) inject energy to the ball, which is crucial when the ball has to be rolled over the fingers with sufficient energy but with a safety margin so that the balls do not overshoot and fall down.

For retaining these critical points while getting rid of unnecessary fluctuations in the trajectory, we recorded impulses in the trajectory before applying the smoothing introduced in the previous subsection. A threshold was set for impulse detection. The trajectory shown in Figure 7 used N=60 Gaussians for smoothing, and the kick threshold was set to 4.5 degrees per cycle. That is, the instants when the desired change in robot joints was larger than 4.5 degrees were recorded as impulses. The desired joint angles of the following and preceding cycles were also recorded. To playback the trajectory on the robot, the smooth analytic representation of the trajectory is resampled at the control frequency of 30Hz. Then the recorded impulse list was superimposed over the smoothed resampled trajectory, preserving the amount of energy injected to the balls during these critical instants.

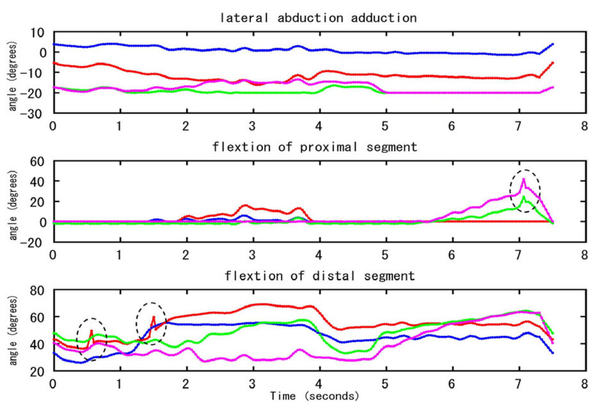

Figure 7. The original robot trajectories were smoothed and the kicks were superimposed as marked with circular dashed curves. The ends of the trajectories were interpolated to meet the beginnings so that the task can be executed as a loop. The shown trajectory produces smooth ball-swapping on the robot.

By experimenting with the kick threshold value we have seen that there were three necessary kicks to complete the task as marked in Figure 7. This trajectory when played back on the robot produced a smooth and very robust ball swapping performance. It looks like it is possible to simplify the trajectory even more as the remaining fluctuations appears to be not critical for the task execution (e.g. see the top panel curves in Figure 7). Note that the end of each curve was linearly interpolated to the beginning so that the task can be executed in a looped fashion.

C. Smoothing and Time Warping with Impulse Preservation

The natural question one asks is whether the ball swapping can be executed faster. One way to test this is to require longer human training in an attempt to obtain faster ball swapping trajectories. The other alternative is to work on the robot side. This can be done by working on the trajectory as we did in the previous subsection, or by implementing a learning algorithm on the robot where it improves some performance measure via practice [11]. Here we present the former approach where we have manipulated the initial trajectory for improving speed.

The straightforward way of scaling the trajectory obtained in the previous section has the disadvantage of washing out the impulses; because in order to play the scaled trajectory on the robot, the trajectory must be resampled, and hence depending on the resampling, the impulses may be lost. The solution is to apply the same trick of the previous section:

1. Find the kicks on the original trajectory
2. Compute a (smooth) analytic representation
3. Speed up: resample the analytic representation
4. Find out the new locations of the kicks based on the resampling rate and superimpose the kicks on (3)

Since the smoothing procedure (step 2) gives an analytic representation of the joint trajectories as $\tilde{\boldsymbol{\theta}}(\mathbf{t}) = \mathbf{k}(\mathbf{t})\mathbf{W}$, it is trivial to speed up the trajectory (say, by a factor of α) via scaling the time with $\tilde{\boldsymbol{\theta}}_{fast}(\mathbf{t}) = \mathbf{k}(\alpha t)\mathbf{W}$ and letting the end time be **T**/α.

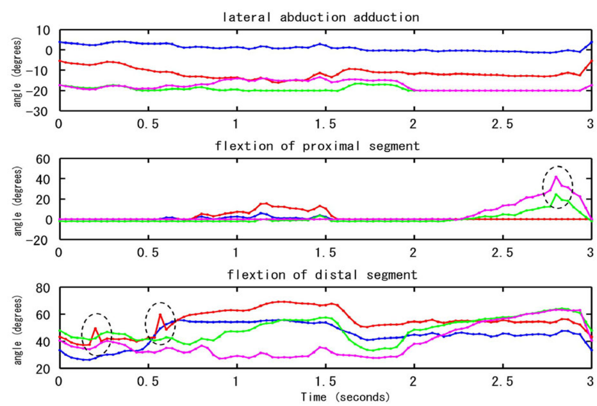

Figure 8. The original trajectory smoothed, sped-up by 2.5 and kicks superimposed. This trajectory produces fast ball-swapping on the robot. The robustness starts to break after this level of speed-up. A movie clip showing the sped-up execution is available at http://www.cns.atr.jp/~erhan/Movies/speedup-320x240.mpg

This modified equation is then used to sample data points at 30Hz completing step 3. Steps 4 is trivially completed by scaling the impulse times recorded in step 1 by 1/α. Figure 8 shows the resulting trajectory, when the speed up factor α is set to 2.5. As the trajectory speeds up the number of data points that are send to the robot per unit time reduces. This puts higher demand on the lower level PD controller. In addition the dynamic interaction of the balls and the robot fingers (e.g. finger-ball contact) becomes non-negligible, eventually breaking the robustness of the execution of ball swapping. With the current scheme of speeding up, our experiments show that we can have speed gains of slightly above 2.5.

VIII. Conclusion

This study indicates that our proposal that the human capacity to learn novel control tasks can be used to synthesize dexterous robot behavior is attainable. In particular, we have demonstrated that two-ball swapping skill can be transferred to a 16-DOF robotic hand through human visuomotor learning. Moreover, we showed that the basic skill obtained through human learning can be improved in terms of smoothness and speed. The analysis of the human generated robot trajectory enabled us to improve the original 7.5 second per swap performance to 3 second per swap, without significant loss of robustness.

VIII. Discussion

In this study, we have focused at the human side of learning but did not address the robot side. We are content that the initial acquired skill can be self-improved via robot learning. This will be much faster than learning from scratch, as in for example, using a reinforcement learning scheme with no prior state or action value function knowledge. In fact the human learning can be used to build up the initial state/action value function to bootstrap further learning. Another possibility is to give certain autonomy and learning capability to the robot while human is learning to control robot behavior. This might yield a very fast mutual learning scheme for both the human and the robot. These robot learning issues are currently being explored in our lab.

The robotic hand in this study can be seen akin to a novel tool that we need to learn to use or interface with, like a computer mouse, that at sometime in the past, we had to spent time to get used to. We think using anthropomorphically similar robots will make it easier for humans to subsume a high DOF robot into the body schema and execute dexterous tasks. It may even be possible to teach humanoid robots to walk like a human within this framework. The sole feedback for the ball swapping task was vision, which was enough for this task; however, for walking a humanoid, at least body orientation must be fed back to the human using specialized training hardware setup that reflects the orientation and angular acceleration of the robot onto the human.

This study, strongly emphasizes that the path to humanoid behavior synthesis can benefit immensely from the body of knowledge and techniques developed in brain and behavioral sciences.

Acknowledgment

This work was supported by JST-ICORP Computational Brain Project. We thank David M. Mount and Sunil Arya for the KD-tree implementation.

References

[1] H. Imamizu, S. Miyauchi, T. Tamada, Y. Sasaki, R. Takino, B. Putz, T. Yoshioka, and M. Kawato, "Human cerebellar activity reflecting an acquired internal model of a new tool," *Nature*, vol. 403, pp. 192-5, 2000.

[2] A. Iriki, M. Tanaka, and Y. Iwamura, "Coding of modified body schema during tool use by macaque postcentral neurones," *Neuroreport*, vol. 7, pp. 2325-30, 1996.

[3] S. Obayashi, T. Suhara, K. Kawabe, T. Okauchi, J. Maeda, Y. Akine, H. Onoe, and A. Iriki, "Functional brain mapping of monkey tool use," *Neuroimage*, vol. 14, pp. 853-61, 2001.

[4] A. Billard and R. Siegwart, "Robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 47, pp. 65-67, 2004.

[5] A. Billard, Y. Epars, S. Calinon, S. Schaal, and G. Cheng, "Discovering optimal imitation strategies," *Robotics and Autonomous Systems*, vol. 47, pp. 69-77, 2004.

[6] C. Breazeal and B. Scassellati, "Robots that imitate humans," *Trends Cogn Sci*, vol. 6, pp. 481-487, 2002.

[7] Y. Demiris and G. Hayes, "Imitation as a dual-route process featuring predictive and learning components: a biologically-plausible computational model," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. Nehaniv, Eds.: MIT Press, 2002.

[8] S. Schaal, "Is imitation learning the route to humanoid robots?," *Trends Cogn Sci*, vol. 3, pp. 233-242, 1999.

[9] S. Schaal, A. Ijspeert, and A. Billard, "Computational approaches to motor learning by imitation," *Philosophical Transaction of the Royal Society of London: Series B, Biological Sciences*, vol. 358, 1431, pp. 537-547, 2003.

[10] A. Ijspeert, J. Nakanishi, and S. Schaal, "Learning attractor landscapes for learning motor primitives," in *Advances in Neural Information Processing Systems*, vol. 15, S. Becker, S. Thrun, and K. Obermayer, Eds. Cambridge, MA:: MIT Press, 2003, pp. 1547-1554.

[11] D. C. Bentivegna, C. G. Atkeson, and G. Cheng, "Learning tasks from observation and practice," *Robotics and Autonomous Systems*, vol. 47, pp. 163-169, 2004.

[12] E. Oztop, L.-H. Lin, M. Kawato, and G. Cheng, "Dexterous Skills Transfer by Extending Human Body Schema to a Robotic Hand," presented at IEEE-RAS International Conference on Humanoid Robots, Genova, Italy, 2006.

[13] C. G. Atkeson and S. Schaal, "Memory-based neural networks for robot learning," *Neurocomputing*, vol. 9, pp. 243-269, 1995.

[14] E. J. Gibson, *Principles of perceptual learning and development*. Englewood Cliffs, NJ: Prentice-Hall, 1969.

[15] A. R. Luft and M. M. Buitrago, "Stages of motor skill learning," *Mol Neurobiol*, vol. 32, pp. 205-16, 2005.

[16] T. Brashers-Krug, R. Shadmehr, and E. Bizzi, "Consolidation in human motor memory," *Nature*, vol. 382, pp. 252-5, 1996.

[17] J. W. Krakauer and R. Shadmehr, "Consolidation of motor memory," *Trends Neurosci*, vol. 29, pp. 58-64, 2006.

[18] R. Shadmehr and H. H. Holcomb, "Neural Correlates of Motor Memory Consolidation," *Science*, vol. 277, pp. 821-825, 1997.

[19] G. Caithness, R. Osu, P. Bays, H. Chase, J. Klassen, M. Kawato, D. M. Wolpert, and J. R. Flanagan, "Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks," *J Neurosci*, vol. 24, pp. 8662-71, 2004.

[20] L. Berthouze, P. Bakker, and Y. Kuniyoshi, "Development of oculo-motor control through robotic imitation.," presented at IEEE/RSJ International Conference on Robotics and Intelligent Systems, Osaka, Japan, 1997.

[21] Y. Kuniyoshi, Y. Yorozu, M. Inaba, and H. Inoue, "From Visuo-Motor Self Learning to Early Imitation - A Neural Architecture for Humanoid Learning," presented at International Conference on Robotics & Automation, Taipei, Taiwan, 2003.

[22] C. G. Atkeson, J. G. Hale, F. E. Pollick, M. Riley, S. Kotosaka, S. Schaal, T. Shibata, G. Tevatia, A. Ude, S. Vijayakumar, and M. Kawato, "Using Humanoid Robots to Study Human Behavior," *IEEE Intelligent Systems*, vol. 15, pp. 46-56, 2000.