16

Sticky Hands

Joshua G. Hale[1] & Frank E. Pollick[2]
[1]*JST ICORP Computational Brain Project, Japan*
[2]*Dept. of Psychology, University of Glasgow, Scotland*

1. Introduction

Sticky Hands is a unique physically-cooperative exercise that was implemented with a full-size humanoid robot. This involved the development of a novel biologically-inspired learning algorithm capable of generalizing observed motions and responding intuitively. Since Sticky Hands has thus far been performed as an exercise for the mutual development of graceful relaxed motion and comfortable physical intimacy between two humans, performing the exercise with a humanoid robot represents a conceptual advance in the role of humanoid robots to that of partners for human self-development.

Engendering a sense of comfortable physical intimacy between human and robot is a valuable achievement: humans must be able to interact naturally with humanoid robots, and appreciate their physical capabilities and requirements in given situations in order to cooperate physically. Humans are attuned to performing such assessments with regard to other humans based on numerous physical and social cues that it will be essential to comprehend and replicate in order to facilitate intuitive interaction.

The chapter expands these issues in more detail, drawing attention to the relevance of established research in motion control for computer graphics and robotics, human motion performance, and perceptual psychology. In particular, attention is paid to the relevance of theories of human motion production, which inspire biomimetic motion performance algorithms, and experimental results shedding light on the acuity of human perception of motion and motion features. We address the following questions: How do we interact naturally, instinctively, intuitively, and effectively with the most adaptable of multipurpose machines, *i.e.*, humanoid robots? How do we make interactions natural, graceful, and aesthetically pleasing? How do we encourage human attribution in the perception of humanoid robots? How can we expand the concept of humanoid robots through novel applications? How can we draw on knowledge already existing in other fields to inspire developments in humanoid robotics?

It is with this perspective, and by way of reply to these inquiries, that we begin this chapter by describing the implementation of the Sticky Hands system, its intuitive learning and generalization system, and hardware. We illustrate how a study of biological processes inspired various aspects of the system, such as the plastic memory of its learning system. We present experimental results and evaluations of the Sticky Hands system.

We conclude the chapter by returning to the main philosophical themes presented in the paper, and describe a fruitful future direction for humanoid robotics research: we conclude that there is a synergistic relationship between neuroscience and humanoid robotics with

Source: Humanoid Robots, New Developments, Book edited by: Armando Carlos de Pina Filho
ISBN 978-3-902613-02-8, pp.582, I-Tech, Vienna, Austria, June 2007

the following four benefits: (i) neuroscience inspires engineering solutions *e.g.*, as we have illustrated, through human motor production and motion perception research (although the breadth of relevant knowledge is much wider than the scope of these two examples); (ii) engineering implementations yield empirical evaluations of neuroscientific theories of human capabilities; (iii) engineering implementations inspire neuroscientific hypotheses; (iv) engineering solutions facilitate new neuroscientific experimental paradigms, *e.g.*, through the recording, analysis and replication of motion for psychovisual experiments.
We believe that exploiting this synergy will yield rapid advances in both fields, and therefore advocate parallel neuroscientific experimentation and engineering development in humanoid robotics. Regarding research building on the Sticky Hands system, experiments involving physical interaction will thus further the development of humanoids capable of interacting with humans, and indeed humans capable of interacting with humanoids.

2. The Sticky Hands game

'Sticky Hands' was drawn from Tai Chi practice, which often includes physical contact exercises with a partner. The game involves forming contact with a partner gently and maintaining it while moving in a previously undetermined pattern. The goal is to develop an ability to perform relaxed and graceful motion, through learning to be sensitive to the forces transmitted via contact and intuitively predict one's partner's movements. When one partner yields the other must push, and vice versa. Prolonged practice reveals the development of intuition and conditioned responses so that the contact may be preserved with a very slight force throughout complex spontaneous sequences of motion. Aspects of the game include becoming comfortable with physical contact, a mutual goal of personal development, and a fulfilling and calming influence. Some people therefore regard it as a form of spiritual development. It is possible for an expert to educate a beginner by encouraging them to perform graceful and rewarding movements, and breaking down the tension in the student's motion.
Our goal was to have a humanoid robot take the role of one partner and play Sticky Hands with a human. In order to rationalise the interaction, we defined a specific variant of the game involving contact using one hand. Partners stand and face each other, raise one hand to meet their partner's and begin making slow circling motions. The path of the contact point may then diverge into a spontaneously developing trajectory as the partners explore their range of physical expression. The robot DB may be seen playing the game in Fig. 1.

Fig. 1. Playing Sticky Hands with humanoid robot DB.

Sticky Hands, being a novel interaction, presents several interesting challenges for the design of an intelligent system, including complex motion planning that mimics the ineffable quality of human intuition. As a basis the robot must be capable of moving while remaining compliant to contact forces from a human, and refining this behaviour, be able to learn from and mimic the motion patterns of humans playing the game. By imitating the types of pattern produced by humans, the robot may reflect creativity and encourage people to explore their own range of motion. To accomplish this while simultaneously experiencing new motion patterns that may develop unpredictably it is necessary to provide a learning algorithm capable of generalising motion it has seen to new conditions, maintaining a suitably evolving internal state.

This work was motivated by a desire to explore physical interaction between humans and robots. This desire grew from the observation that the arising problem for the design of appropriate intelligent systems is how to communicate and cooperate with people (Takeda et al.1997). Using the principle that physical interaction is a familiar and reliable method for most humans, Sticky Hands broaches the area of physical cooperation and the subtle but significant issue of communication (Coppin et al. 2000) through physical movement (Adams et al. 2000; Hikiji 2000). Our work therefore considers human imitation (Scassellati 2000), which we consider to be of particular relevance for the future -as computer science and robotics develop it becomes clear that humans and robots will cooperate with a wide range of tasks. Moreover, we explored the use of a humanoid robot as a playmate facilitating a human’s self-development. In this context the robot assumes a new social role involving a physically intimate and cooperative interaction. We can also hope that through the interaction, people will be encouraged to attribute the robot an anthropomorphic identity rather than considering it as a mechanical entity. Such a shift of perspective heralds ways of making human and robot interactions more natural.

Humanoid robotics embodies a certain fascination with creating a mechanical entity analogous to our human selves. There are of course many other valid motivations for creating humanoids (Bergener et al. 1997), not least among which is the nature of our environment - being highly adapted to human sensory and motor capabilities it begs for artificial agents with analogous capabilities that can usefully coexist in our own living and working spaces. Having an anthropomorphic shape and organic motion also makes working with such robots aesthetically and emotionally more pleasing. The production of human-like, and emotionally expressive styles of movement are of particular relevance. The Sticky Hands interaction embodies human-like motion and autonomy. The target of maintaining minimal contact force is tangible. The creative, anticipatory aspect however, is enhanced by initiative. The challenges posed by these problems motivated the development of a highly generalized learning algorithm, and a theoretical investigation of expressive motion styles.

We continue this section with a system overview describing the relationship between robot control and learning in the Sticky Hands system. We then outline the robot control issues, describing additional sensing technology and explaining how we achieved hand placement compliant to external forces. We then discuss the learning algorithm which observes trajectories of the hand throughout interaction with a human and predicts the development of a current trajectory.

2.1. System overview

The control system for robotic Sticky Hands was treated as three components which are shown in Fig. 2. The *robot motor controller* is responsible for positioning the hand, and obeys a planned trajectory supplied by the learning algorithm. It also estimates the contact force between the human's and robot's hands, adjusting the trajectory plan to compensate for contact force discrepancies. The actual hand position data was smoothed to remove noise, and sent back to the learning algorithm.

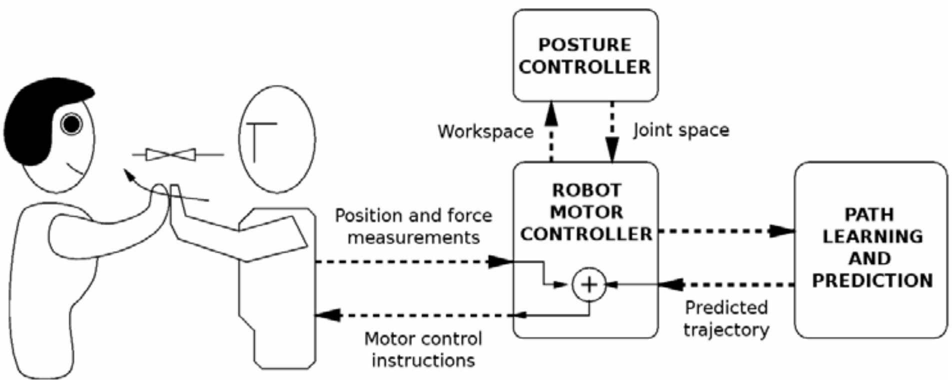

Fig. 2. System breakdown.

The *path learning and prediction* component is a learning algorithm that outputs a predicted hand trajectory and processes the actual observed trajectories supplied by the motor controller. The learning algorithm observes the evolution of the hand trajectory continuously. It learns motion patterns, and generalises them to predict future developments in the hand trajectory. The input and output are both sequences of position vectors. The robot controller makes use of the *posture controller*. The posture controller uses a straightforward inverse kinematics routine to generate joint configurations satisfying Cartesian hand placement targets.

2.2 Robot control

The Sticky Hands exercise was performed by a 30 DOF SARCOS anthropomorphic robot (Atkeson et al. 2000) that may be seen in Fig. 3. Each joint is powered hydraulically, and has angle and load sensors. Joints are servoed independently by the low-level controller using torques proportional to the angular offset between the measured and target angles, and negatively proportional to the angular velocity at each joint. This yields proportional gains/spring-damper control at each joint, where the torque at a given joint is calculated as:

$$\tau = k_s(\theta_t - \theta) + k_d(\dot{\theta}_t - \dot{\theta}) \tag{1}$$

θ_t and $\dot{\theta}_t$ are the target angle and angular velocity (usually $\dot{\theta}_t = 0$), θ and $\dot{\theta}$ are the current angle and angular velocity. k_s and k_d are the spring stiffness and damping parameters respectively.

Oscillations caused by this proportional gains controller were avoided by employing an inverse dynamics algorithm to estimate the torques necessary to hold a position. Since the robot is anchored off the ground by its pelvis, standing and balancing did not constitute problems. The Sticky Hands exercise involves only one hand and the chain of joints from the anchor point to the robot's hand encompasses 10 DOFs. The chain is kinematically redundant, so an iterative inverse kinematics algorithm was used (Tevatia & Schaal 2000). The 20 unused DOFs were resolved according to a default posture.

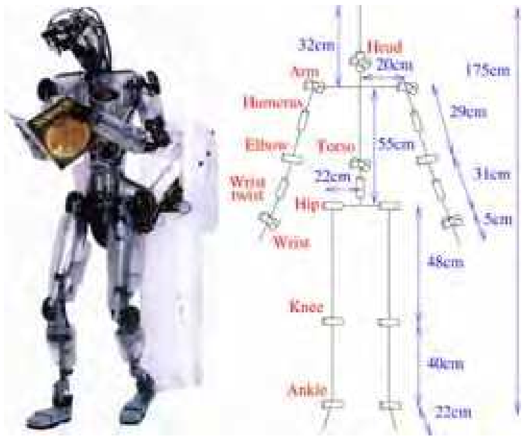

Fig. 3. DB kinematics.

The robot was required to balance a force applied by the human player against its hand. Trajectories were performed while incorporating any adjustments needed to balance the contact force, thus making the hand *actively compliant* to changes of the contact force. The simplest method of determining the contact force is direct measurement, *i.e.*, by means of a force transducer between the robot's and human's hands. Fig. 4 shows the attachment of a force transducer between the robot's hand and a polystyrene hemisphere intended to facilitate an ergonomic surface for the human to contact. The transducer measured forces in the *X*, *Y* and *Z* directions. The target position of the hand was translated to compensate if the contact force exceeded a 5N threshold. In order to establish a balance point against the force applied by the human, we subtracted a small quantity in the *Z* (forward-backward) direction from the measured force prior to thresholding. We also implemented a method of responding to the contact force using only the sensors internal to the SARCOS robot, *i.e.*, joint load and angle. In both cases we assumed that the interactions were not so forceful as to necessitate more than a constant adjustment in hand position, *i.e.*, continuous pressure would result in a yielding movement with a constant velocity.

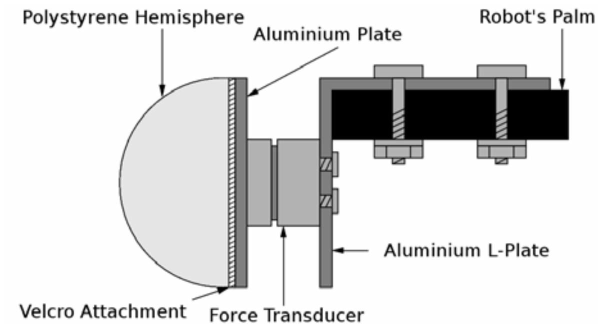

Fig. 4. Force transducer attachment.

It is possible to calculate the externally applied force using the internal sensors of DB. This involves measuring the joint angles, estimating the torques necessary to hold the position using inverse dynamics, and subtracting the measured loads. Any discrepancy should be torques due to external loads other than gravity, because gravity is the only external force accounted for in the inverse dynamics calculation. The method relies on having a very accurate inverse dynamics model, and unfortunately inaccuracies in the dynamic model and load sensors necessitated a large contact force to facilitate suitable thresholding of the results.

We therefore measured the positional offset between target and actual hand positions instead, assuming any large *positional* discrepancy was caused by external non-gravity forces. We were able to threshold by 2cm (yielding an effective force threshold of about 12N). When the low-level controller was required to follow a suggested trajectory, the hand was directed to 5cm further forward from the suggested position. The human was assumed to supply a force sufficient to maintain the hand in the actual target position specified by the suggested trajectory. This method did not require readings from the joint load sensors to respond to external forces because physical interactions were registered purely through the position of the robot's hand -which requires only joint angle readings to compute. If the human reduced or increased the contact force the robot's hand would move beyond the 2cm threshold and the suggested trajectory consequently adjusted to compensate. While the contact force was entirely in the Z direction, perturbations in the X and Y as well as Z directions were monitored, and accommodated by translating the target position when the 2cm threshold was reached. This indirect kinematic method facilitated the use of a significantly lighter contact force than with inverse dynamics based force estimation.

Any reasonable positioning and compliance strategy is in fact compatible with the high-level Sticky Hands control system. Ideally, the force threshold is minimal, since the threshold determines the minimum contact force. The larger the contact force the less relaxed the human's motion may be. The way the redundancy in robot hand placement is resolved does not significantly affect the contact force but on the other hand may have an effect on the appearance of the robot's motion. We discuss this point later. We also compare this kinematic hand positioning technique with the force transducer method and present traces of the forces measured during interaction with a human.

The trajectories supplied by the learning algorithm were described using piecewise linear splines. The robot controller ran at 500Hz and the learning algorithm ran at 10Hz. The sequence of predictions output by the learning algorithm was interpreted as the advancing end-point of a spline. The knots of this spline were translated to compensate for any discrepancy in the contact force detected by the motor controller. This translation was accomplished smoothly, in order to prevent the hand from jerking in response to contact forces: after translating the knots, a negative translation vector was initialized to bring the knots back to their original position. This negative translation was gradually decayed to zero during each cycle of the motor controller. The sum of the translated knots and the negative translation vector thus interpolated from the original knot positions and the translated knot positions.

2.3 'Prototype set' learning algorithm

3D point samples describing the robot's hand trajectory were fed as input into the learning algorithm. A vector predicting the progression of the trajectory was output in return for each sample supplied. The learning algorithm fulfilled the following properties, which are required by the nature of the Sticky Hands exercise.

- **Generalise** observed trajectories for prediction of similar new trajectories with different orientation, scale, curvature, position or velocity
- **Extrapolate** properties of a new trajectory for prediction in the absence of similar observed trajectories
- Maintain a **fluid internal state** in order to cope with the evolving nature of trajectories through continuous update, replacement and 'forgetting' of recorded information
- **Handle branch points** where similar observed trajectories diverge
- Tolerate **noise** causing inaccuracy in position samples
- Facilitate a **parameterisable time bound,** ensuring real time operation
- Facilitate a **parameterisable memory bound,** in order to fully exploit the host architecture

We refer to the structure used to record the instantaneous properties of the input trajectory as a 'prototype'. This paradigm may be compared to the work of Stokes et al. (1999) who presented a method for identifying cyclic patterns and their significance in *space-line* samples. Our process focuses rather on the immediate instant of a trajectory. Salient features are recorded for efficient retrieval but no internal classifications of higher level structures such as cycles are made. This is the essence of the generalisation and branch point handling properties of our algorithm since the recorded properties of any instant of an observed trajectory may be used to predict the development of any new trajectory. Moreover, no information about correlations between trajectories is maintained in an explicit form by the prediction process.

The 'prototype' is defined mathematically below and the creation of prototypes from raw geometrical information is presented. The utilisation of prototypes for *prediction* and *extrapolation* is also demonstrated. Then the issue of how to select the most appropriate prototype for predicting a given trajectory from a memory bank of prototypes is addressed. The memory bank is maintained according to a reinforcement principle designed to ensure an efficient use of memory by 'forgetting' prototypes that are not necessary. This search procedure requires a distance metric between prototypes and has been optimised. The reader may find it useful to refer to Fig. 5 throughout this prototype learning section.

2.3.1 Prediction using prototypes

Given a sequence of input position samples, $\{p_k : k \in \mathrm{N}\}$. The prototype P_i corresponding to p_i is defined using p_{i-1}, p_i and p_{i+1}:

$$P_i = (p_i, v_i, a_i, T_i) \tag{2}$$

$$v_i = p_{i+1} - p_i \in \mathrm{R}^3 \tag{3}$$

$$a_i = \frac{|p_{i+1} - p_i|}{|p_i - p_{i-1}|} \in \mathrm{R} \tag{4}$$

$$T_i = \left(\cos\left(\frac{\theta_i}{2}\right), \sin\left(\frac{\theta_i}{2}\right)(p_{i+1} - p_i) \times (p_i - p_{i-1}) \right) \tag{5}$$

$$\theta_i = \cos^{-1}\left(\frac{(p_i - p_{i-1}).(p_{i+1} - p_i)}{|p_i - p_{i-1}||p_{i+1} - p_i|} \right) \in \mathrm{R} \tag{6}$$

v_i is the velocity out of p_i (scaled by an arbitrary factor) and a_i is a scalar indicating the magnitude of the acceleration. The *direction* of the acceleration is deducible from T_i, which is a quaternion describing the change in direction between v_i and v_{i-1} as a rotation through their mutually orthogonal axis.

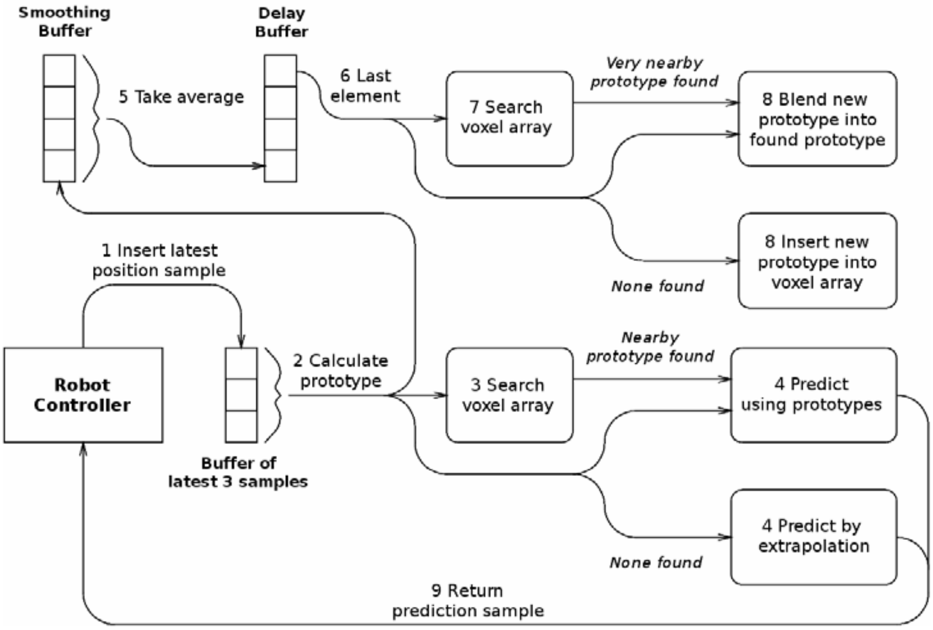

Fig. 5. Datapath in the learning algorithm (arrows) and execution sequence (numbers).

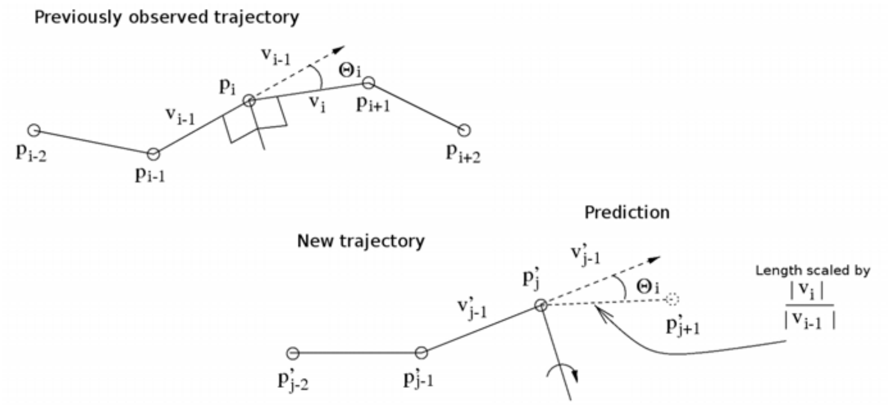

Fig. 6. Trajectory prediction using a prototype.

The progression of a trajectory $\{p'_k : k \in \mathrm{N}\}$ at a given instant may be predicted using a prototype. Suppose that for a particular trajectory sample p'_j, it is known that P_i

corresponds best to p'_j , then $p'_j + a_i T_i (p'_j - p'_{j-1})$ is an estimate for p'_{j+1} . Pre-multiplication of a 3-vector by T_i denotes quaternion rotation in the usual way. This formula applies the bend and acceleration occurring at p_i to predict the position of p'_j . We also linearly blend the position of p_i into the prediction, and the magnitude of the velocity so that p'_j combines the actual position and velocity of p_i with the prediction duplicating the bending and accelerating characteristics of p_i (see Fig. 6):

$$p'_{j+1} = p'_j + s_j T_i \cdot \frac{p'_j - p'_{j-1}}{| p'_j - p'_{j-1} |} + g_p (p_i - p'_j) \tag{7}$$

$$s_j = (1 - g_v) a_i \, | p'_j - p'_{j-1} | + g_v \, | v_i | \tag{8}$$

g_p and g_v are blending ratios used to manage the extent to which predictions are entirely general, or repeat previously observed trajectories, *i.e.*, how much the robot wants to repeat what it has observed. We chose values of g_p and g_v in the range [0.1, 0.001] through empirical estimation. g_p describes the tendency of predictions to gravitate spatially towards recorded motions, and g_v has the corresponding effect on velocity.

In the absence of a corresponding prototype we can calculate P'_{j-1} , and use it to estimate p'_{j+1} , thus extrapolating the current characteristics of the trajectory. Repeated extrapolations lie in a single plane determined by p_{i-2}, p_{i-1} and p_i , and maintain the trajectory curvature (rotation in the plane) measured at p'_j . We must set $g_p = 0$ since positional blending makes no sense when extrapolating, and would cause the trajectory to slow to a halt, *i.e.*, the prediction should be based on an extrapolation of the immediate velocity and turning of the trajectory and not averaged with its current position since there is no established trajectory to gravitate towards.

2.3.2 Storage and retrieval

Ideally, when predicting p'_{j+1} , an observed trajectory with similar characteristics to those at p'_j is available. Typically a large set of recorded prototypes is available, and it is necessary to find the closest matching prototype P_i or confirm that no suitably similar prototype exists. The prototype P'_{j-1} which is generated from the current trajectory can be used as a basis for identifying similar prototypes corresponding to similar, previously observed trajectories. We define a distance metric relating prototypes in order to characterise the closest match.

$$d(P_i, P_j) = 1 - \cos(\theta') + \frac{| p_i - p_j |}{M_p} \tag{9}$$

Where,

$$\theta \in [-M_a, M_a] \Rightarrow \theta' = \theta \frac{\pi}{2M_a}$$
$$\theta \notin [-M_a, M_a] \Rightarrow \theta' = \pi \tag{10}$$

$$\theta = \cos^{-1} \frac{v_i . v_j}{|v_i||v_j|} \tag{11}$$

M_a and M_p define the maximum angular and positional differences such that $d(P_i, P_j)$ may be one or less. Prototypes within this bound are considered similar enough to form a basis for a prediction, *i.e.*, if $d(P_i, P_j)$ is greater than 1 for all *i* then no suitably similar prototype exists. The metric compares the position of two prototypes, and the direction of their velocities. Two prototypes are closest if they describe a trajectory traveling in the same direction, in the same place. In practice, the values of 15cm and $\pi/4$ radians for M_p and M_a respectively were found to be appropriate. -A trajectory with exactly the same direction as the developing trajectory constitutes a match up to a displacement of 15cm, a trajectory with no displacement constitutes a match up to an angular discrepancy of $\pi/4$ radians, and within those thresholds there is some leeway between the two characteristics. The threshold values must be large enough to permit some generalisation of observed trajectories, but not so large that totally unrelated motions are considered suitable for prediction when extrapolation would be more appropriate.

The absolute velocity, and bending characteristics are not compared in the metric. Predictions are therefore general with respect to the path leading a trajectory to a certain position with a certain direction and velocity, so branching points are not problematic. Also the speed at which an observed trajectory was performed does not affect the way it can be generalised to new trajectories. This applies equally to the current trajectory and previously observed trajectories.

When seeking a prototype we might naïvely compare *all* recorded prototypes with P'_{j-1} to find the closest. If none exist within a distance of 1 we use P'_{j-1} itself to extrapolate as above. Needless to say however, it would be computationally over-burdensome to compare P'_{j-1} with *all* the recorded prototypes. To optimise this search procedure we defined a voxel array to store the prototypes. The array encompassed a cuboid enclosing the reachable space of the robot, partitioning it into a $50 \times 50 \times 50$ array of cuboid voxels indexed by three integer coordinates. The storage requirement of the empty array was 0.5MB. New prototypes were placed in a list attached to the voxel containing their positional component p_i. Given P_{j-1} we only needed to consider prototypes stored in voxels within a distance of M_p from p_{j-1} since prototypes in any other voxels would definitely exceed the maximum distance according to the metric. Besides limiting the total number of candidate prototypes, the voxel array also facilitated an optimal ordering for considering sets of prototypes. The voxels were considered in an expanding sphere about p_j. A list of integer-triple voxel index offsets was presorted and used to quickly identify voxels close to a given centre voxel ordered by minimum distance to the centre voxel. The list contained voxels up to a minimum distance of M_p. This ensures an optimal search of the voxel array since the search may terminate as soon as we encounter a voxel that is too far away to contain a prototype with a closer minimum distance than any already found. It also permits the search to be cut short if time is unavailable. In this case the search terminates optimally since the voxels most likely to contain a match are considered first. This facilitates the parameterisable time bound since the prototype search is by far the dominant time expense of the learning algorithm.

2.3.3 Creation and maintenance

Prototypes were continually created based on the stream of input position samples describing the observed trajectory. It was possible to create a new prototype for each new sample, which we placed in a cyclic buffer. For each new sample we extracted the average prototype of the buffer to reduce sampling noise. A buffer of 5 elements was sufficient. The averaged prototypes were shunted through a delay buffer, before being added to the voxel array. This prevented prototypes describing a current trajectory from being selected to predict its development (extrapolation) when other prototypes were available. The delay buffer contained 50 elements, and the learning algorithm was iterated at 10Hz so that new prototypes were delayed by 5 seconds.

Rather than recording every prototype we limited the total number stored by averaging certain prototypes. This ensures the voxel array does not become clogged up and slow, and reduces the memory requirement. Therefore before inserting a new prototype into the voxel array we first searched the array for a similar prototype. If none was found we added the new prototype, otherwise we blended it with the existing one. We therefore associated a count of the number of blends applied to each prototype to facilitate correct averaging with new prototypes. In fact we performed a non-linear averaging that capped the weight of the existing values, allowing the prototypes to tend towards newly evolved motion patterns within a limited number of demonstrations. Suppose P_a incorporates n blended prototypes, then a subsequent blending with P_b will yield:

$$P_a' = P_a \frac{D(n)-1}{D(n)} + P_b \frac{1}{D(n)} \tag{12}$$

$$D(n) = 1 + A_M - \frac{A_M}{1 + nA_G} \tag{13}$$

$A_M/(1 + nA_M)$ defines the maximum weight for the old values, and A_G determines how quickly it is reached. Values of 10 and 0.1 for A_M and A_G respectively were found to be suitable. This makes the averaging process linear as usual for small values but ensures the contribution of the new prototype is worth at least 1/11th.

We facilitated an upper bound on the storage requirements using a deletion indexing strategy for removing certain prototypes. An integer clock was maintained, and incremented every time a sample was processed. New prototypes were stamped with a deletion index set in the future. A list of the currently stored prototypes sorted by deletion index was maintained, and if the storage bounds were reached the first element of the list was removed and the corresponding prototype deleted. The list was stored as a *heap* (Cormen et al.) since this data structure permits fast $O(\log(\text{num elements}))$ insertion, deletion and repositioning. We manipulated the deletion indices to mirror the reinforcement aspect of human memory. A function $R(n)$ defined the period for which a prototype reinforced n times should be retained (n is equivalent to the blending count). Each time a prototype was blended with a new one we calculated the retention period, added the current clock and re-sorted the prototype index. $R(n)$ increases exponentially up to a maximum asymptote.

$$R(n) = D_M - \frac{D_M}{1 + D_G n^{D_P}} \tag{14}$$

D_M gives the maximum asymptote. D_G and D_P determine the rate of increase. Values of 20000, 0.05 and 2 were suitable for D_M, D_G and D_P respectively. The initial reinforcement thus extended a prototype's retention by 2 minutes, and subsequent reinforcements roughly doubled this period up to a maximum of about half an hour (the algorithm was iterated at 10Hz).

3. Results

The initial state and state after playing Sticky Hands with a human partner are shown in Fig. 7. Each prototype is plotted according to its position data. The two data sets are each viewed from two directions and the units (in this and subsequent figures) are millimeters. The *X, Y* & *Z* axes are positive in the robot's left, up and forward directions respectively. The point *(0,0,0)* corresponds to the robot's sacrum. The robot icons are intended to illustrate orientation only, and not scale. Each point represents a unique prototype stored in the motion predictor's memory, although as discussed each prototype may represent an amalgamation of several trajectory samples. The trajectory of the hand loosely corresponds to the spacing of prototypes but not exactly because sometimes new prototypes are blended with old prototypes according to the similarities between each's position and velocity vectors.

The initial state was loaded as a default. It was originally built by teaching the robot to perform an approximate circle 10cm in radius and centred in front of the left elbow joint (when the arm is relaxed) in the frontal plane about 30cm in front of the robot. The prototype positions were measured at the robot's left hand, which was used to play the game and was in contact with the human's right hand throughout the interaction. The changes in the trajectory mostly occur gradually as human and robot slowly and cooperatively develop cycling motions. Once learned, the robot can switch between any of its previously performed trajectories, and generalise them to interpret new trajectories.

The compliant positioning system, and its compatibility with motions planned by the prediction algorithm was assessed by comparing the Sticky Hands controller with a 'positionable hand' controller that simply maintains a fixed target for the hand in a compliant manner so that a person may reposition the hand.

Fig. 8 shows a force/position trace where the width of the line is linearly proportional to the magnitude of the force vector (measured in all 3 dimensions), and Table 1 shows corresponding statistics. Force measurements were averaged over a one minute period of interaction, but also presented are 'complied forces', averaging the force measurements over only the periods when the measured forces exceeded the compliance threshold. From these results it is clear that using the force transducer yielded significantly softer compliance in all cases. Likewise the 'positionable hand' task yielded slightly softer compliance because the robot did not attempt to blend its own trajectory goals with those imposed by the human.

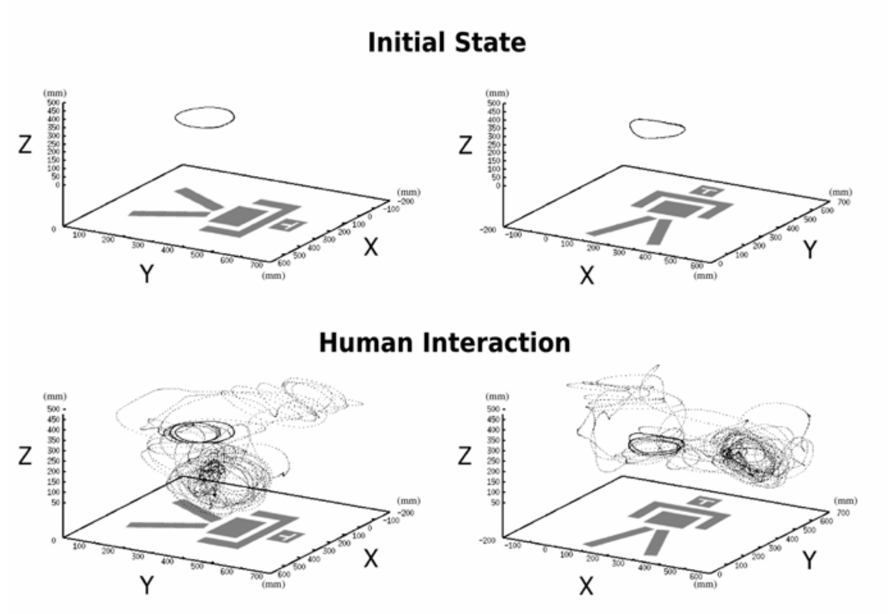

Fig. 7. Prototype state corresponding to a sample interaction.

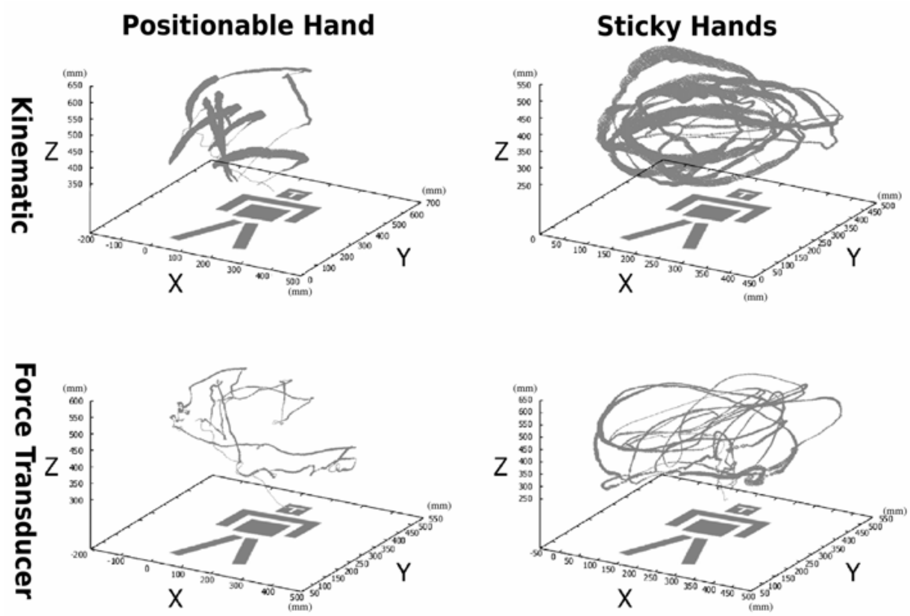

Fig. 8. Force measured during 'positionable hand' and Sticky Hands tasks.

Task	Contact force (N)		Complied forces (N)	
	Mean	Var.	Mean	Var.
Force Transducer Sticky Hands	4.50	4.83	5.72	4.49
Force Transducer ‘Positionable Hand’	1.75	2.18	3.23	2.36
Kinematically Compliant Sticky Hands	11.86	10.73	13.15	10.73
Kinematically Compliant ‘Positionable Hand’	8.90	10.38	12.93	11.40

Table 1 Forces experienced during ‘positionable hand’ and Sticky Hands tasks.

Examining a sequence of interaction between the robot and human reveals many of the learning system’s properties. An example sequence during which the robot used the kinematic compliance technique is shown in Fig. 9. The motion is in a clockwise direction, defined by progress along the path in the a-b-c direction, and was the first motion in this elliptical pattern observed by the prediction system. The ‘Compliant Adjustments’ graph shows the path of the robot’s hand, and is marked with thicker lines at points where the compliance threshold was exceeded. *i.e.*, points where the prediction algorithm was mistaken about the motion the human would perform. The ‘Target Trajectory’ graph shows in lighter ink the target sought by the robot’s hand along with in darker ink the path of the robot’s hand. The target is offset in the Z (forwards) direction in order to bring about a contact force against the human’s hand. At point *(a)* there is a kink in the actual hand trajectory, a cusp in the target trajectory, and the beginning of a period during which the robot experiences a significant force from the human. This kink is caused by the prediction algorithm’s expectation that the trajectory will follow previously observed patterns that have curved away in the opposite direction, the compliance maintaining robot controller adjusts the hand position to attempt to balance the contact force until the curvature of the developing trajectory is sufficient to extrapolate its shape and the target trajectory well estimates the path performed by the human. At point *(b)* however, the human compels the robot to perform an elliptical shape that does not extrapolate the curvature of the trajectory thus far. At this point the target trajectory overshoots the actual trajectory due to its extrapolation. Once again there is a period of significant force experienced against the robot’s hand and the trajectory is modified by the compliance routine. At point *(c)* we observe that, based on the prototypes recorded during the previous ellipse, the prediction algorithm correctly anticipates a similar elliptical trajectory offset positionally and at a somewhat different angle.

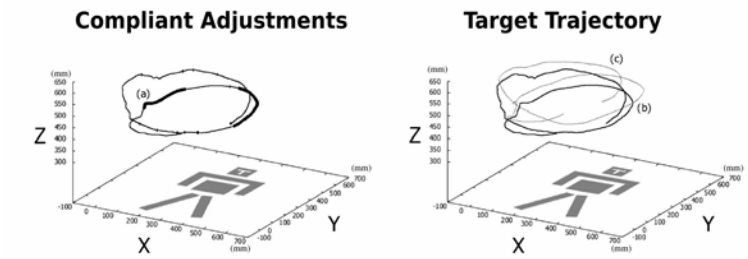

Fig. 9 Example interaction showing target trajectory and compliance activation

4. Discussion

We proposed the ‘Sticky Hands’ game as a novel interaction between human and robot. The game was implemented by combining a robot controller process and a learning algorithm with a novel internal representation. The learning algorithm handles branching trajectories implicitly without the need for segmentation analysis because the approach is not pattern based. It is possible to bound the response time and memory consumption of the learning algorithm arbitrarily within the capabilities of the host architecture. This may be achieved trivially by restricting the number of prototypes examined or stored. The ethos of our motion system may be contrasted with the work of Williamson (1996) who produced motion controllers based on positional primitives. A small number of postures were interpolated to produce target joint angles and hence joint torques according to proportional gains. Williamson's work advocated the concept of ``behaviours or skills as coarsely parameterised atoms by which more complex tasks can be successfully performed''. Corresponding approaches have also been proposed in the computer animation literature, such as the motion verbs and adverbs of Rose et al. (1998). Williamson's system is elegant, providing a neatly bounded workspace, but unfortunately it was not suitable for our needs due to the requirements of a continuous interaction incorporating more precise positioning of the robot's hand.

By implementing Sticky Hands, we were able to facilitate physically intimate interactions with the humanoid robot. This enables the robot to assume the role of playmate and partner assisting in a human's self-development. Only minimal sensor input was required for the low-level motor controller. Only torque and joint position sensors were required, and these may be expected as standard on most humanoid robots. With the addition of a hand mounted force transducer the force results were also obtained. Our work may be viewed as a novel communication mechanism that accords with the idea that an autonomous humanoid robot should accept command input and maintain behavioral goals at the same level as sensory input (Bergener et al. 1997). Regarding the issue of human instruction however, the system demonstrates that the blending of internal goals with sensed input can yield complex behaviors that demonstrate a degree of initiative. Other contrasting approaches (Scassellati 1999) have achieved robust behaviors that emphasize the utility of human instruction in the design of reinforcement functions or progress estimators.

The design ethos of the Sticky Hands system reflects a faith in the synergistic relationship between humanoid robotics and neuroscience. The project embodies the benefits of cross-fertilized research in several ways. With reference to the introduction, it may be seen that (i) neuroscientific and biological processes have informed and inspired the development of the system, *e.g.,* through the plastic memory component of the learning algorithm, and the control system's "intuitive" behaviour which blends experience with immediate sensory information as discussed further below; (ii) by implementing a system that incorporates motion based social cues, the relevance of such cues has been revealed in terms of human reactions to the robot. Also, by demonstrating that a dispersed representation of motion is sufficient to yield motion learning and generalization, the effectiveness of solutions that do not attempt to analyze nor segment observed motion has been confirmed; (iii) technology developed in order to implement Sticky Hands has revealed processes that could plausibly be used by the brain for solving motion tasks, *e.g.,* the effectiveness of the system for blending motion targets with external forces to yield a compromise between the motion modeled internally and external influences suggests that humans might be capable of performing learned motion patterns according to a consistent underlying model subject to forceful external influences that might significantly alter the final motion; (iv) the Sticky Hands system is in itself a valuable tool for research since it provides an engaging cooperative interaction between a human and a humanoid robot. The robot 's behaviour

may be modulated in various ways to investigate for example the effect of less compliant motion, different physical cues, or path planning according to one of various theories of human motion production.
The relationship between the engineering and computational aspect of Sticky Hands and the neuroscientific aspect is thus profound. This discussion is continued in the following sections which consider Sticky Hands in the context of relevant neuroscientific fields: human motion production, perception, and the attribution of characteristics such as naturalness and affect. The discussion is focused on interaction with humans, human motion, and lastly style and affect.

4.1 Interacting with humans

The Sticky Hand task requires two partners to coordinate their movements. This type of coordination is not unlike that required by an individual controlling an action using both their arms. However, for such bimanual coordination there are direct links between the two sides of the brain controlling each hand. Though surprisingly, even when these links are severed in a relatively rare surgical intervention known as callosotomy, well-learned bimanual processes appear to be remarkably unaffected (Franz, Waldie & Smith, 2000). This is consistent with what we see from experienced practitioners of Tai Chi who perform Sticky Hands: that experience with the task and sensory feedback are sufficient to provide graceful performance. It is a reasonable speculation that the crucial aspect of experience lays in the ability to predict which movements are likely to occur next, and possibly even what sensory experience would result from the actions possible from a given position.
A comparison of this high level description with the implementation that we used in the Sticky Hands task is revealing. The robot's experience is limited to the previous interaction between human and robot and sensory information is limited to either the kinematics of the arm and possibly also force information. Clearly the interaction was smoother when more sensory information was available and this is not entirely unexpected. However, the ability of the robot to perform the task competently with a very minimum of stored movements is impressive. One possibility worth considering is that this success might have been due to a fortunate matching between humans' expectations of how the game should start and the ellipse that the robot began with. This matching between human expectations and robot capabilities is a crucial question that is at the heart of many studies of human-robot interaction.
There are several levels of possible matching between robot and human in this Sticky Hands task. One of these, as just mentioned is that the basic expectations of the range of motion are matched. Another might be that the smoothness of the robot motion matches that of the human and that any geometric regularities of motion are matched. For instance it is known that speed and curvature are inversely proportional for drawing movements (Lacquaniti et al. 1983) and thus it might be interesting in further studies to examine the effect of this factor in more detail. A final factor in the relationship between human and robot is the possibility of social interactions. Our results here are anecdotal, but illustrative of the fact that secondary actions will likely be interpreted in a social context if one is available. One early test version of the interaction had the robot move its head from looking forward to looking towards its hand whenever the next prototype could not be found. From the standpoint of informing the current state of the program this was useful. However, there was one consequence of this head movement that likely was exacerbated by the fact that it was the more mischievous actions of the human partner that would confuse the robot. This lead the

robot head motion to fixate visually on its own hand, which by coincidence was where most human partners were also looking, leading to a form of mutual gaze between human and robot. This gestural interaction yielded variable reports from the human players as either a sign of confusion or disapproval by the robot.

This effect is illustrative of the larger significance of subtle cues embodied by human motion that may be replicated by humanoid robots. Such actions or characteristics of motion may have important consequences for the interpretation of the movements by humans. The breadth of knowledge regarding these factors further underlines their value. There is much research describing how humans produce and perceive movements and many techniques for producing convincing motion in the literature of computer animation. For example, there is a strong duality between dynamics based computer animation and robotics (Yamane & Nakamura 2000). Computer animation provides a rich source of techniques for generating (Witkin & Kass 1988; Cohen 1992; Ngo & Marks 1993; Li et al. 1994; Rose et al. 1996; Gleicher 1997) and manipulating (Hodgins & Pollard 1997) dynamically correct motion, simulating biomechanical properties of the human body (Komura & Shinagawa 1997) and adjusting motions to display affect or achieve new goals (Bruderlin & Williams 1995; Yamane & Nakamura 2000).

4.2 Human motion

Although the technical means for creating movements that appear natural and express affect, skill, *etc.* are fundamental, it is important to consider the production and visual perception of *human* movement. The study of human motor control for instance holds the potential to reveal techniques that improve the replication of human-like motion. A key factor is the *representation* of movement. Interactions between humans and humanoids may improve if both have similar representations of movement. For example, in the current scenario the goal is for the human and robot to achieve a smooth and graceful trajectory. There are various objective ways to express smoothness. It can be anticipated that if both the humanoid and human shared the same representation of smoothness then the two actors may converge more quickly to a graceful path.

The visual perception of human movement likewise holds the potential to improve the quality of human-robot interactions. The aspects of movement that are crucial for interpreting the motion correctly may be isolated according to an analysis of the features of motion to which humans are sensitive. For example, movement may be regarded as a complicated spatiotemporal pattern, but the recognition of particular styles of movement might rely on a few isolated spatial or temporal characteristics of the movement. Knowledge of human motor control and the visual perception of human movement could thus beneficially influence the design of humanoid movements.

Several results from human motor control and motor psychophysics inform our understanding of natural human movements. It is generally understood several factors contribute to the smoothness of human arm movements. These include the low-pass filter characteristics of the musculoskeletal system itself, and the planning of motion according to some criteria reflecting smoothness. The motivation for such criteria could include minimizing the wear and tear on the musculoskeletal system, minimizing the overall muscular effort, and maximizing the compliance of motions. Plausible criteria that have been suggested include the minimization of jerk, *i.e.,* the derivative of acceleration (Flash & Hogan 1985), minimizing the torque change (Uno et al. 1989), the motor-command change (Kawato 1992), or signal dependent error (Harris & Wolpert 1998).

There are other consistent properties of human motion besides smoothness that have been observed. For example, that the endpoint trajectory of the hand behaves like a concatenation of piecewise planar segments (Soechting & Terzuolo 1987a; Soechting & Terzuolo 1987b). Also, the movement speed is related to its geometry in terms of curvature and torsion. Specifically, it has been reported that for planar segments velocity is inversely proportional to curvature raised to the 1/3rd power, and that for non-planar segments the velocity is inversely proportional to the 1/3rd power of curvature multiplied by 1/6th power of torsion (Lacquaniti et al. 1983; Viviani & Stucchi 1992; Pollick & Sapiro 1996; Pollick et al. 1997; Handzel & Flash, 1999). Extensive psychological experiments of the paths negotiated by human-humanoid dyads could inform which principles of human motor control are appropriate for describing human-humanoid cooperative behaviours.

4.3 Style and affect

Recent results examining the visual recognition of human movement are also of relevance with regard to the performance of motion embodying human-like *styles*. By considering the relationship between movement kinematics and style recognition, it has been revealed that recognition can be enhanced by exaggerating temporal (Hill & Pollick 2000), spatial (Pollick et al. 2001a), and spatiotemporal (Giese & Poggio 2000; Giese & Lappe 2002) characteristics of motion. The inference of style from human movement (Pollick et al. 2001b) further supports the notion that style may be specified at a kinematic level. The kinematics of motion may thus be used to constrain the design of humanoid motion.

However, the meaningful kinematic characteristics of motion may rely on dynamic properties in a way that can be exploited for control purposes. The brief literature review on human motor control and visual perception of human movement above provides a starting point for the design of interactive behaviours with humanoid robots. The points addressed focus on the motion of the robot and may be viewed as dealing with the problem in a bottom up fashion. In order to make progress in developing natural and affective motion it is necessary to determine whether or not a given motion embodies these characteristics effectively. However, it is possible that cognitive factors, such as expectancies and top down influences might dominate interactions between humans and humanoids, *e.g.,* the humanoid could produce a natural movement with affect but the motion could still be misinterpreted if there is an expectation that the robot would not move naturally or display affect.

5. Conclusion

Having described the Sticky Hands project: it's origin, hardware and software implementation, biological inspiration, empirical evaluation, theoretical considerations and implications, and having broadened the later issues with a comprehensive discussion, we now return to the enquiries set forth in the introduction.

The Sticky Hands project itself demonstrates a natural interaction which has been accomplished effectively –the fact that the objectives of the interaction are in some aspects open-ended creates leeway in the range of acceptable behaviours but also imposes complex high-level planning requirements. Again, while these may be regarded as peculiar to the Sticky Hands game they also reflect the breadth of problems that must be tackled for advanced interactions with humans. The system demonstrates through analysis of human motion, and cooperation how motion can be rendered naturally, gracefully and aesthetically. These characteristics are both key objectives in Sticky Hands interaction, and as we have indicated in the discussion also have broader implications for the interpretation, quality and effectiveness of interactions with humans in general for which the attribution of human qualities such as emotion engender an expectation of the natural social cues that improve the effectiveness of cooperative behaviour through implicit communication.

We have drawn considerable knowledge and inspiration from the fields of computer graphics, motion perception and human motion performance. The benefit that the latter two fields offer for humanoid robotics reveal an aspect a larger relationship between humanoid robotics and neuroscience. There is a synergistic relationship between the two fields that offers mutual inspiration, experimental validation, and the development of new experimental paradigms to both fields. We conclude that exploring the depth of this relationship is a fruitful direction for future research in humanoid robotics.

8. References

Adams, B.; Breazeal, C.; Brooks, R.A. & Scassellati, B. (2000). Humanoids Robots: A New Kind of Tool. *IEEE Intelligent Systems*, 25-31, July/August

Atkeson, C.G.; Hale, J.G.; Kawato, M.; Kotosaka, S.; Pollick, F.E.; Riley, M.; Schaal, S.; Shibata, T.; Tevatia, G.; Ude A. & Vijayakumar, S. (2000). Using humanoid robots to study human behavior. *IEEE Intelligent Systems*, 15, pp46-56

Bergener, T.; Bruckhoff, C.; Dahm, P.; Janben, H.; Joublin, F. & Menzner, R. (1997). Arnold: An Anthropomorphic Autonomous Robot for Human Environments. *Proc. Selbstorganisation von Adaptivem Verhalten (SOAVE 97)*, 23-24 Sept., Technische Universitt Ilmenau

Bruderlin, A. & Williams, L. (1995). Motion Signal Processing. *Proc. SIGGRAPH 95, Computer Graphics Proceedings*, Annual Conference Series, pp97-104

Cohen, M.F. (1992). Interactive Spacetime Control for Animation. *Proc. SIGGRAPH 92, Computer Graphics Proceedings*, Annual Conference Series, pp293-302

Coppin P.; Pell, R.; Wagner, M.D.; Hayes, J.R.; Li, J.; Hall, L. ; Fischer, K.D.; Hirschfield & Whittaker, W.L. (2000). EventScope: Amplifying Human Knowledge and Experience via Intelligent Robotic Systems and Information Interaction. *IEEE International Workshop on Robot-Human Interaction*, Osaka, Japan

Cormen, T.H.; Leiserson, C.E. & Rivest, R.L. Introduction To Algorithms. McGraw-Hill, ISBN 0-07-013143-0

Flash, T. & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, 5, pp1688-1703

Giese, M.A. & Poggio, T. (2000). Morphable models for the analysis and synthesis of complex motion patterns. *International Journal of Computer Vision*, 38, 1, pp59-73

Giese, M.A. & Lappe, M. (2002). Perception of generalization fields for the recognition of biological motion. *Vision Research*, 42, pp1847-1858

Gleicher, M. (1997). Motion Editing with Spacetime Constraints. *Proc. 1997 Symposium on Interactive 3D Graphics*

Handzel, A. A. & Flash, T. (1999). Geometric methods in the study of human motor control. *Cognitive Studies* 6, 309-321

Harris, C.M. & Wolpert, D.M. (1998). Signal-dependent noise determines motor planning. *Nature* 394, pp780-784

Hikiji, H. (2000). Hand-Shaped Force Interface for Human-Cooperative Mobile Robot. *Proceedings of the 2000 IEICE General Conference*, A-15-22, pp300

Hill, H. & Pollick, F.E. (2000). Exaggerating temporal differences enhances recognition of individuals from point light displays. *Psychological Science*, 11, 3, pp223-228

Hodgins, J.K. & Pollard, N.S. (1997). Adapting Simulated Behaviors For New Characters. *Proc. SIGGRAPH 97, Computer Graphics Proceedings*, Annual Conference Series, pp153-162

Kawato, M. (1992). Optimization and learning in neural networks for formation and control of coordinated movement. In: *Attention and performance*, Meyer, D. and Kornblum, S. (Eds.), XIV, MIT Press, Cambridge, MA, pp821-849

Komura, T. & Shinagawa, Y. (1997). A Muscle-based Feed-forward controller for the Human Body. *Computer Graphics forum* 16(3), pp165-176

Lacquaniti, F.; Terzuolo, C.A. & Viviani, P. (1983). The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54, pp115-130

Li, Z.; Gortler, S.J. & Cohen, M.F. (1994). Hierarchical Spacetime Control. *Proc. SIGGRAPH 94, Computer Graphics Proceedings*, Annual Conference Series, pp35-42

Ngo, J.T. & Marks, J. (1993). Spacetime Constraints Revisited. *Proc. SIGGRAPH 93, Computer Graphics Proceedings*, Annual Conference Series, pp343-350

Pollick, F.E. & Sapiro, G. (1996). Constant affine velocity predicts the 1/3 power law of planar motion perception and generation. *Vision Research,* 37, pp347-353

Pollick, F.E.; Flash, T.; Giblin, P.J. & Sapiro, G. (1997). Three-dimensional movements at constant affine velocity. *Society for Neuroscience Abstracts*, 23, 2, pp2237

Pollick, F.E.; Fidopiastis, C.M. & Braden, V. (2001a). Recognizing the style of spatially exaggerated tennis serves. *Perception*, 30, pp323-338

Pollick, F.E.; Paterson, H.; Bruderlin, A. & Sanford, A.J. (2001b). Perceiving affect from arm movement. *Cognition*, 82, B51-B61

Rose, C.; Guenter, B.; Bodenheimer, B. & Cohen, M.F. (1996). Efficient Generation of Motion Transitions using Spacetime Constraints. *Proc. SIGGRAPH 96, Computer Graphics Proceedings*, Annual Conference Series, pp147-154

Rose, C.; Bodenheimer, B. & Cohen, M.F. (1998). Verbs and Adverbs: Multidimensional Motion Interpolation. *IEEE Computer Graphics & Applications,* 18(5)

Scassellati, B. (1999). Knowing What to Imitate and Knowing When You Succeed. *Proc. of AISB Symposium on Imitation in Animals and Artifacts*, Edinburgh, Scotland

Scassellati, B. (2000). Investigating models of social development using a humanoid robot. In: *Biorobotics*, Webb, B. and Consi, T. (Eds.), MIT Press, Cambridge, MA

Soechting, J.F. & Terzuolo, C.A. (1987a). Organization of arm movements. Motion is segmented. *Neuroscience*, 23, pp39-51

Soechting, J.F. & Terzuolo, C.A. (1987b). Organization of arm movements in three-dimensional space. Wrist motion is piecewise planar. *Neuroscience*, 23, pp53-61

Stokes, V.P.; Lanshammar, H. & Thorstensson, A. (1999). Dominant Pattern Extraction from 3-D Kinematic Data. *IEEE Transactions on Biomedical Engineering* 46(1)

Takeda H.; Kobayashi N.; Matsubara Y. & Nishida, T. (1997). Towards Ubiquitous Human-Robot Interaction. *Proc. of IJCAI Workshop on Intelligent Multimodal Systems*, Nagoya Congress Centre, Nagoya, Japan

Tevatia, G. & Schaal, S. (2000). Inverse kinematics for humanoid robots. *IEEE International Conference on Robotics and Automation*, San Francisco, CA

Uno, Y.; Kawato, M. & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, 61, pp89-101

Viviani, P. & Stucchi, N. (1992). Biological movements look uniform: Evidence of motor-perceptual interactions. *Journal of Experimental Psychology: Human Perception and Performance*, 18, pp602-623

Williamson, M.M. (1996). Postural primitives: Interactive Behavior for a Humanoid Robot Arm. *Proc. of SAB 96*, Cape Cod, MA, USA

Witkin, A. & Kass, M. (1988). Spacetime Constraints. *Proc. SIGGRAPH 88, Computer Graphics Proceedings*, Annual Conference Series, pp159-168

Yamane, K. & Nakamura, Y. (2000). Dynamics Filter: Towards Real-Time and Interactive Motion Generator for Human Figures. *Proc. WIRE 2000*, 27-34, Carnegie Mellon University, Pittsburgh, PA

Humanoid Robots: New Developments
Edited by Armando Carlos de Pina Filho

ISBN 978-3-902613-00-4
Hard cover, 582 pages
Publisher I-Tech Education and Publishing
Published online 01, June, 2007
Published in print edition June, 2007

For many years, the human being has been trying, in all ways, to recreate the complex mechanisms that form the human body. Such task is extremely complicated and the results are not totally satisfactory. However, with increasing technological advances based on theoretical and experimental researches, man gets, in a way, to copy or to imitate some systems of the human body. These researches not only intended to create humanoid robots, great part of them constituting autonomous systems, but also, in some way, to offer a higher knowledge of the systems that form the human body, objectifying possible applications in the technology of rehabilitation of human beings, gathering in a whole studies related not only to Robotics, but also to Biomechanics, Biomimmetics, Cybernetics, among other areas. This book presents a series of researches inspired by this ideal, carried through by various researchers worldwide, looking for to analyze and to discuss diverse subjects related to humanoid robots. The presented contributions explore aspects about robotic hands, learning, language, vision and locomotion.