

International Journal of Humanoid Robotics
 © World Scientific Publishing Company

Accurate Task-Space Tracking for Humanoids with Modeling Errors using Iterative Learning Control

Pranav A. Bhounsule

*Dept. of Mechanical Engineering,
 The University of Texas San Antonio,
 One UTSA Circle, San Antonio, TX 78249, USA.
 pranav.bhounsule@utsa.edu*

Katsu Yamane

*Disney Research,
 4720 Forbes Avenue, Lower Level, Suite 110,
 Pittsburgh, PA 15213, USA.
 kyamane@disneyresearch.com*

Precise task-space tracking with manipulator-type systems requires an accurate kinematic model. In contrast to traditional manipulators, sometimes it is difficult to obtain an accurate kinematic model of humanoid robots due to complex structure and link flexibility. Also, prolonged use of the robot will lead to some parts wearing out or being replaced with a slightly different alignment, thus throwing off the initial calibration. Therefore, there is a need to develop a control algorithm that can compensate for the modeling errors and quickly retune itself, if needed, taking into account the controller bandwidth limitations and high dimensionality of the system. In this paper, we develop an iterative learning control algorithm that can work with existing inverse kinematics solvers to refine the joint-level control commands to enable precise tracking in the task space. We demonstrate the efficacy of the algorithm on a theme-park type humanoid doing a drawing task, serving drink in a glass, and serving a drink placed on a tray without spilling. The iterative learning control algorithm is able to reduce the tracking error by at least two orders of magnitude in less than 20 trials.

Keywords: Iterative Learning Control, Task-space tracking, Humanoids, Zero-Reference Model, Constraint Optimization

1. Introduction

Many applications of manipulator-type systems involve precise task-space control. Examples include industrial manipulators that do welding and personal humanoid robots that folds clothes. Industrial manipulators rely on good CAD models and high-gain servo controllers to accomplish precise end-effector motions. However, the techniques that work well for industrial manipulators may not transfer to humanoid robots for the following reasons:

- (1) Humanoid robots are made light weight due to weight and size constraints leading to flexible links. Thus, traditional CAD models which are based on

rigid body assumptions are not valid. Additional modeling terms are needed to account for link flexibility, which can be hard.

- (2) Humanoid robots tend to have small actuators for safety reasons and/or may have a low bandwidth controller. This makes it hard to implement a precise servo.
- (3) Humanoid robots have large number of joints and long limbs, especially the human-size robots. Thus, small parameters errors in the CAD model can lead to big errors at the end-effectors.
- (4) When the humanoid robots are used long-term (e.g., personal robots), some parts may wear out or be replaced with a slightly different alignment. Hence the original CAD model is not valid anymore.

Thus, for accurate task-space control of humanoid robots, one needs a method that can compensate for the modeling errors by modifying the joint-level control commands and to make up for part wear and/or replacement. We present an iterative learning control algorithm that can address the above issues.

In this paper, we show that a combination of constrained optimization and iterative learning control using data from motion capture (MoCap) system can enable high fidelity task-space tracking even with strict joint limits. First, we create a rigid body based kinematics model and fit the parameters by minimizing the error between model predicted end-effector pose and measured pose using the MoCap system. When we invert the obtained kinematic model for task space control, the accuracy is limited because of unmodelled effects, specifically joint flexibility and sensor noise, and position tracking errors due to limited control bandwidth. We therefore employ iterative learning control to improve the tracking performance. The iterative learning control algorithm modifies the desired end-effector motion based on end-effector tracking errors. We show that the learning algorithm is able to provide accurate tracking on three tasks on a theme-park type humanoid: (i) drawing the figure eight, (ii) serving a drink without spilling, and (iii) serving a drink placed on a tray.

2. Background and Related Work

The main issue with task-space control is the lack of accurate kinematic models due to reasons mentioned earlier. Traditional feedback control methods such as Proportional-Integral-Derivative Control ¹ are the preferred methods to correct for modeling errors because they have a simple structure and can be relatively easy to hand-tune. However, they are not preferred when the plant is subject to unexpected disturbances. In this case, it is common to have an adaptive controller that modifies the control parameters to make up for the varying loads ^{2,3}. However, most feedback control techniques rely on setting high gains which necessitates the use of high bandwidth feedback control, typically 500 Hz or more. In our case, the control bandwidth of 120 Hz limits us to relatively low gains.

Learning-based method have also been used to do task-space control. One ap-

proach is to directly build a inverse kinematics model using the experimental data^{4,5}. The biggest issue with this approach is that the inverse mapping is not unique⁶. To overcome the multiple solution nature of the inverse mapping, Oyama et al.⁷ used a multiple neural networks to represent the inverse kinematic solutions locally in different regions of the state space. These individual networks are called experts. Next, another neural network, called the gating network, is used to choose an expert to obtain the inverse kinematics solution. One of the problem with the above method is that the construction of the gating network becomes difficult in high dimensions.

Iterative Learning Control (ILC) can correct for modeling errors to enable high fidelity tracking. In its simplest form, ILC modifies the control command in every iteration in the following way: the command at trial i , is the sum of the command in trial $i - 1$, and a control gain times the tracking error in trial $i - 1$ ⁸. Because the tracking errors are reduced iteratively at every trial, the control gains can be kept small.

Traditionally, in ILC, the learning is at the joint level⁹. However, it is quite straightforward to extend ILC to task level using the appropriate mapping from the task-space to the joint space. For example, Arimoto et al.^{10,11} used the linearized mapping, i.e., the Jacobian, to map the incremental change in position from the task level to the joint level and showed the efficacy of their algorithm on a four link manipulator in simulation. In our ILC algorithm, we use the non-linear map from task-space to joint space and show the efficacy of the algorithm experimentally on a humanoid robot. Specifically, we evaluate the non-linear map by using an inverse kinematics solver which finds a solution within joint limits. The main improvement over Arimoto's algorithm is that our algorithm is able to handle joint limits (see Bhounsule et al.¹² for more details). This paper is an extension of the paper we previously published¹³ in the following ways: we give extensive details about the kinematic modeling including experimental results and provide an extensive evaluation of the iterative learning control algorithm on a number of tasks on the humanoid robot.

3. Hardware

3.1. Humanoid robot

We use a 37 degree of freedom, hydraulically powered, fixed-base humanoid robot shown in Fig. 1. Each joint has either a rotary potentiometer or a linear variable differential transformer to sense the joint position. There are two levels of control. At the lowest level, there is a single processor per joint. The processor runs a 1 KHz control loop that does high gain position control using individual position data from joint sensor, velocity from differentiated and filtered position data from joint sensor and the force sensor in the valve. The gains on the lowest level controller are pre-set and cannot be changed at runtime. At the highest level, there is a single computer that communicates with all the low level processors at 120 Hz, sending desired

joint position commands and receiving measured positions. The highest level serves as our interface to the robot, i.e., the input is a position command and the joint position is the measured variable, both of which occur at a rate of 120 Hz.

3.2. Marker-based motion capture

We use OptiTrack motion capture (MoCap) system along with their software suite ARENA¹⁴. We use 8 cameras to track the 3 passive retro-reflective markers attached at each of the end-effectors: the head, and the two hands. Using the position data from the 3 markers, the ARENA software computes the position and the orientation of each of the end-effectors.

The motion capture system is first used to create a kinematic model of the robot and subsequently used for iterative learning control to improve on the kinematic model for accurate task space tracking. In all experiments, the position and the orientation of the end-effectors from the MoCap system and the robot's joint position in potentiometer ticks are recorded simultaneously at 120 Hz.

4. Kinematic model for humanoid

4.1. Joint model

We assume a linear relation between the joint position in radians or meters and the measured joint position signal in counts from the potentiometers or linear variable differential transformers. Our model is

$$q_i = K_i(q_{i,\text{counts}} - q_{i,\text{counts}}^0), \quad (1)$$

where i denotes the joint number; K_i is the gain in radians per count for revolute joint or meters per count for prismatic joint; q_i is the angle in radians (denoted θ) or linear position in meters (denoted as s); $q_{i,\text{counts}}$ is the measured angle or linear position in counts, a digital signal from 0 to 4095, as we use a 12 bit encoder; $q_{i,\text{counts}}^0$ is the bias for the angle or for the position.

4.2. Kinematic model

We use the zero-reference kinematic modeling^{15,16,17,18}. In the zero reference modeling approach the zero position (the reference configuration in which all joint position have a value equal to zero) can be arbitrarily fixed¹⁹, unlike the more popular Denavit-Hartenberg modeling approach²⁰ where the pose depends on the robot geometry. Because we use the zero reference kinematic modeling approach, we need to have a fixed zero reference configuration, and is chosen as shown in Fig. 1 (a). Next, we choose the unit vector that defines the direction of motion of the joint axis. By defining a point through which the axis passes, we obtain the location of the joint axis¹⁹.

The homogenous transformation T_i defines the motion of a point on the link attached to the joint i ($i = 1, 2, \dots, 26, A, B, \dots, E$) with respect to the previous joint.

We define the homogenous transformation, \mathbf{T}_i , for revolute and prismatic joints next.

Revolute joint

$$\mathbf{T}_i^{\text{revolute}} = \begin{bmatrix} \mathbf{R} & (\mathbf{I} - \mathbf{R})\mathbf{r}^T \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (2)$$

where \mathbf{I} is the 3×3 identity matrix; $\mathbf{r} = [x_i, y_i, z_i]$ is 3×1 vector that gives location of the revolute joint with respect to previous joint in the local reference frame; and \mathbf{R} is a 3×3 rotation matrix and is shown next.

$$\mathbf{R} = \begin{bmatrix} u_x^2 v\theta + c\theta & u_x u_y v\theta - u_z s\theta & u_x u_z v\theta + u_y s\theta \\ u_x u_y v\theta + u_z s\theta & u_y^2 v\theta + c\theta & u_y u_z v\theta - u_x s\theta \\ u_x u_z v\theta - u_y s\theta & u_y u_z v\theta + u_x s\theta & u_z^2 v\theta + c\theta \end{bmatrix},$$

where θ is the angle turned by the joint in radians; $\mathbf{u} = [u_x, u_y, u_z]$ defines the joint axis; and $s\theta = \sin \theta$, $c\theta = \cos \theta$, $v\theta = 1 - \cos \theta$. To obtain the transformation given by Equation (2) we do the following: first translate the point by a distance $-\mathbf{r}$ to pass through the center of the revolute joint; then rotate by an angle θ using the Rodrigues equation; and finally translate by \mathbf{r} to restore it to the original position.

The transform, $\mathbf{T}_i^{\text{revolute}}$, has 5 unknowns; two in \mathbf{u} , two in \mathbf{r} , and one in θ . We obtain the third component of \mathbf{u} using the identity, $u_x^2 + u_y^2 + u_z^2 = 1$. By definition, the zero-reference model fixes one component of \mathbf{r} . Hence, there are just two unknowns in the position vector \mathbf{r} .

Prismatic joint

$$\mathbf{T}_i^{\text{prismatic}} = \begin{bmatrix} 1 & 0 & 0 & s u_x \\ 0 & 1 & 0 & s u_y \\ 0 & 0 & 1 & s u_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (3)$$

where we define \mathbf{u} as in the previous section and s defines the linear motion of the joint in meters.

The transform, $\mathbf{T}_i^{\text{prismatic}}$, has 3 unknowns; two in \mathbf{u} , and one in s . We obtain the third component of \mathbf{u} using the identity $u_x^2 + u_y^2 + u_z^2 = 1$.

The transformation \mathbf{T}_e^i maps the i -th end-effector frame with the zero reference pose. The zero reference pose is shown in Fig. 1 (b).

$$\mathbf{T}_e^i = \begin{bmatrix} \mathbf{R}_e^i & \mathbf{r}_e^i \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

Here, $\mathbf{r}_e^i = [x_e^i, y_e^i, z_e^i]$ is 3×1 vector defines the location of a marker with respect to the previous joint in the local reference frame, and we obtain \mathbf{R}_e^i from Euler angles

$\phi_x^i, \phi_y^i, \phi_z^i$, using the $x-y-z$ convention of Euler angles (also called Bryant angles).

$$\mathbf{R}_e^i = \mathbf{R}_x^i \mathbf{R}_y^i \mathbf{R}_z^i.$$

The transformation given by Equation (4) has 6 parameters in all: three for the position (x_e^i, y_e^i, z_e^i) and three for the orientation ($\phi_x^i, \phi_y^i, \phi_z^i$).

Like Equation (4), the transformation T_r that maps the robot reference frame to the camera reference frame in the zero reference pose is

$$\mathbf{T}_r = \begin{bmatrix} \mathbf{R}_r & \mathbf{r}_r \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (5)$$

From Fig. 1 (b), we see that the humanoid has a closed loop structure from the waist down (grey thick lines that constitute robot legs) and a tree structure from the waist up (thin multi-color lines that constitute the robot torso, the head, and the hands). We break the loop joint at joint 4. Now we have four serial chains: three from each of the end-effectors (the head, the left hand, and the right hand) to joint 1, and one from joint 4 to joint 3. Next, we show how to combine these serial chains to give the kinematic model and the end-effector pose.

If \mathbf{T}_c^i is the transformation that maps the motion of the i -th end-effector with respect to camera frame then

$$\mathbf{T}_c^{\text{head}} = \mathbf{T}_r \mathbf{T}_1 \mathbf{T}_2 \mathbf{T}_A \mathbf{T}_B \mathbf{T}_4 \mathbf{T}_5 \mathbf{T}_6 \mathbf{T}_7 \mathbf{T}_8 \mathbf{T}_9 \mathbf{T}_{10} \mathbf{T}_e^{\text{head}}, \quad (6)$$

$$\mathbf{T}_c^{\text{left-hand}} = \mathbf{T}_r \mathbf{T}_1 \mathbf{T}_2 \mathbf{T}_A \mathbf{T}_B \mathbf{T}_4 \mathbf{T}_5 \mathbf{T}_6 \mathbf{T}_7 \mathbf{T}_{11} \mathbf{T}_{12} \mathbf{T}_{13} \mathbf{T}_{14} \mathbf{T}_{15} \mathbf{T}_{16} \mathbf{T}_{17} \mathbf{T}_{18} \mathbf{T}_e^{\text{left-hand}}, \quad (7)$$

$$\mathbf{T}_c^{\text{right-hand}} = \mathbf{T}_r \mathbf{T}_1 \mathbf{T}_2 \mathbf{T}_A \mathbf{T}_B \mathbf{T}_4 \mathbf{T}_5 \mathbf{T}_6 \mathbf{T}_7 \mathbf{T}_{19} \mathbf{T}_{20} \mathbf{T}_{21} \mathbf{T}_{22} \mathbf{T}_{23} \mathbf{T}_{24} \mathbf{T}_{25} \mathbf{T}_{26} \mathbf{T}_e^{\text{right-hand}}. \quad (8)$$

Finally, to ensure that the loop joint constraint closes at joint 3, we define the transformation $\mathbf{T}_r^{\text{loop}}$ as

$$\mathbf{T}_r^{\text{loop}} = \mathbf{T}_4 \mathbf{T}_C \mathbf{T}_D \mathbf{T}_E \mathbf{T}_3. \quad (9)$$

The transformation $\mathbf{T}_r^{\text{loop}}$ gives the position and the orientation of the end of the loop joint, that is, at joint 3. For the loop to close, we need to satisfy 5 constraints (1 less than a maximum of 6, because of the single degree of freedom at joint 3). The 5 constraints are enforced by choosing appropriate values for the 5 angles at joints, **A**, **B**, **C**, **D**, and **E**. These 5 angles are neither sensed nor actuated on the robot.

We use the kinematics modeling utility in the rigid body dynamics simulator SDFast²² to model the robot kinematics.

4.3. Initial model parameters

The initial model parameter estimation is done as follows. We first set the robot in the zero-reference pose as shown in the schematic diagram in Fig. 1 (a). The positive direction for the joint axis corresponds to the direction that the joint should move

for the raw joint position measurement system to give an increasing value. Setting up the joint axis direction this way ensures that the gains (K_i 's) are positive. We measure link lengths using a ruler. We obtain the bias terms, $q_{i,\text{counts}}^0$ from the measured position in counts in the reference pose. A rough estimate of the gain K_i , for each joint, is obtained as follows. We move the joint whose gain we want to estimate by a random amount. We measure the change in the angle in radians using a protractor, and also in tics using the potentiometer on the joint. We can estimate K_i by dividing the change in angle in radians to the change in angle in tics.

4.4. Parameter Identification

For kinematic modeling, we need a rich data set. We drive each joint using a sinusoidal trajectory that covers the full range of joint motion. The frequencies of individual joints are chosen to be non-multiples of each other. This ensures that we were not getting repeating motion. The total data collection time is about 30 minutes. However, for analysis purposes, we down-sample the data by a factor of 240 to get about 750 unique poses. A kinematic model is specified and parameters are fit using non-linear least squares. The model has limited accuracy due to unmodeled effects such as sensor noise and link deflection due to load. Thus, when the kinematic model is inverted for task-space control there are errors. We therefore use an iterative learning control algorithm to modify the control command based on error between the predicted end-effector position and orientation using the inverse kinematics model and the measurements from the motion capture system to enable accurate task space tracking (see Sec. 5.2).

The parameter identification step improves on the initial estimates obtained earlier using the data collected from MoCap. This is done by updating the model parameters based on minimization of the squared prediction error using the parameter optimization software called SNOPT²³, a non-linear constraint optimization software based on sequential quadratic programming.

Our parameter optimization problem is as follows: For a given set of poses specified by the joint position data in counts, we find parameters for the model given by Equation (1) and Equations (6) to (9), that best explains the measured end-effector position and orientation data from all three end-effectors, and across all the poses.

We use the a single cost function consisting of sum of end-effector position and

orientation error

$$\text{Cost} = \underbrace{\frac{\sum_{i=1}^{N_{\text{head}}} (\Delta r_i^2 + w_i \Delta \theta_i^2)}{N_{\text{head}}}}_{\text{Cost-head}} + \underbrace{\frac{\sum_{i=1}^{N_{\text{left-hand}}} (\Delta r_i^2 + w_i \Delta \theta_i^2)}{N_{\text{left-hand}}}}_{\text{Cost-left-hand}} + \underbrace{\frac{\sum_{i=1}^{N_{\text{right-hand}}} (\Delta r_i^2 + w_i \Delta \theta_i^2)}{N_{\text{right-hand}}}}_{\text{Cost-right-hand}}, \quad (10)$$

$$(\Delta r_i)^2 = (\Delta x_i)^2 + (\Delta y_i)^2 + (\Delta z_i)^2, \quad (11)$$

$$(\Delta \theta_i)^2 = (\Delta \theta_i^x)^2 + (\Delta \theta_i^y)^2 + (\Delta \theta_i^z)^2, \quad (12)$$

where $\Delta x_i = x_i^{\text{model}} - x_i^{\text{expt}}$ denotes the difference between values from the kinematic model (obtained from appropriate transform in Equation (6), Equation (7) or Equation (8)) and the measurement for the x-coordinate of one of the end-effector for the i th pose and so on. Similarly, $\Delta \theta_i^x = \theta_i^{x,\text{model}} - \theta_i^{x,\text{expt}}$ denotes the difference between values from the kinematic model and the measurement for the Euler angle corresponding to rotation along x axis for the i th pose and so on. N_{head} , $N_{\text{left-hand}}$, $N_{\text{right-hand}}$ denote the number of observations for the head, the left-hand, and the right-hand respectively. The weight, w_i denotes a weight on the angular position with respect to the linear position and was set to 1 for the parameter identification.

We now describe the optimization parameters. We have a total of 178 parameters given in the list below. Item 1 gives the parameters for the joint model (see Equation (1)) and the rest of the items are for the kinematic model (see Equation (2) to Equation (9)).

- (1) We have 2 parameters, K_i and $q_{i,\text{counts}}^0$, for each joint. As there are 26 joints, we have a total of 52 joint level parameters.
- (2) We specify the position and the orientation of the i -th rigid body formed by the markers with respect to the end-effector using the transformation \mathbf{T}_e^i and given by Equation (4). This transformation has 6 parameters; three for the position and three for the orientation. For the three end-effectors, there are a total of 18 parameters
- (3) We specify the position and the orientation of the robot's reference frame with respect to the camera frame using the transformation \mathbf{T}_r and given by Equation (5). This transformation has 6 parameters; three for position and three for orientation.
- (4) The tree structure from joint 5 or Torso Twist and up (thin multi-colored lines in Fig. 1 (b)) has 20 revolute joints and 2 sliding joints. Each revolute joint adds 4 parameters (one less than the usual 5 for revolute joints because we know the joint position in radians from the joint model, see Equation (2)); and each sliding joint adds 2 parameters (one less than the normal 3 for prismatic joints because we know the joint position in meters from joint model, see Equation

- (3)). Thus we have a total of 92 ($20 \times 4 + 2 \times 6$) parameters for the tree structure.
- (5) The closed loop structure in the lower body (shown as thick grey line in Fig. 1) has 9 revolute joints but only 4 revolute joints are sensed and actuated. In each actuated joint, there are 2 numbers that describe the joint axis. Since there are 4 actuated joints, there are a total of 8 parameters. The only length parameters in the loop structure are the height and width of the loop. Thus, the loop structure has a total of 10 parameters.

4.5. Results for kinematic modeling

Next, we present results for the parameter identification problem described in Sec. 4.4. Table 1 gives the root mean square error for the position and the orientation for the head, the left hand, and the right hand. The average position accuracy is 1 cm and the average orientation accuracy is 1.84 degrees. For a 1.8 m tall robot, an accuracy of 1 cm corresponds to a percentage error of only 0.6 %, which is reasonably small, but not small enough for accurate task space tracking, as we will see in the next section.

The residual errors for the head are lower than the left hand and the right hand, indicating a better fit for the head. Both hands have similar residual errors. We believe that the higher residual error in the hands as compared to the head is because of two reasons: link deflection due to gravity loading and sensor noise. Since the hands have a bigger mass and larger moment arm than the head, the hands would have a bigger end-effector deflection than the head. Also, since the hands have more degrees of freedom and therefore, more joint sensors, the effect of sensor noise is more dominant for the hand than the head. We could have improved the fit further by adding a model that accounts for link deflection due to gravity loading²⁴, and a model accounting for sensor accuracy²⁵. However, we decided to forego this because we can easily improve the tracking accuracy for joint and task space control by using learning or feedback control algorithms.

Fig. 2 shows a histogram of the position and the orientation error. The errors are normally distributed with a mean of zero, except for the position error in gravity direction, Δz_i , which has a mean around 0.5 cm. We speculate that this error bias is due to gravity loading which we have not accounted for in our model.

Consistency check: We do three more non-linear least squares calculations for individual end-effectors separately using the individual cost terms defined in the Equation (13). We compare these fits with each other and with the least squares fit for all the end-effectors taken together described by all terms in the Equation (13). To check the consistency we compare some of the common parameters in the four optimizations, namely the pin alignment and the link lengths in the lower body. The maximum variation between the four optimizations for the pin axis is 4 degrees and for the link length is 4 cm, primarily because of parameter redundancy (e.g., the height of the closed loop and the height of the mid-body). However, since the

variation we obtain across different optimizations are within 2%, we conclude that the solutions are consistent.

5. Task-level control

We use the inverse of the identified kinematic model for task-level control but found it not accurate enough for precise tracking. We therefore used an iterative learning control algorithm to improve the tracking performance and described in this section.

5.1. Inverse Kinematics

We need an inverse kinematics solver to map the desired end-effector motion to joint space for the low level position servo. Sometimes it is possible to find a closed form inverse kinematics solution (e.g., ²⁶), but this is not always the case. In such instances, a numerical optimization provides a straight-forward and generalizable method of finding an inverse kinematics solution, especially for redundant robots such as humanoids ²⁷. Specifically, by choosing a suitable cost function one can bias the solution to use certain joints over the other ones.

We use the optimization software SNOPT ²³ again, to develop an inverse kinematics solver. The problem here is to find the joint angles as a function of time, $\boldsymbol{\theta}(t)$, to minimize the cost function g , subject to end-effector constraints \mathbf{h}_1 , the lower joint limits \mathbf{h}_2 , and upper joint limits \mathbf{h}_3 . We define these functions next:

$$g(\boldsymbol{\theta}(t)) = \sum_{i=1}^{n_{\text{dof}}} w_i \left(\theta_i(t) - \theta_i^{\text{ref}}(t) \right)^2 + \sum_{j=1}^{\hat{n}_{\text{eff}}} W_j \left(X_{j,\text{des}} - f_j(\boldsymbol{\theta}(t)) \right)^2, \quad (13)$$

$$\mathbf{h}_1(\boldsymbol{\theta}(t)) = \mathbf{X}_{i,\text{des}}(t) - \mathbf{f}_i(\boldsymbol{\theta}(t)) \leq |\epsilon|, \quad \text{where } i = 1, 2, \dots, n_{\text{eff}}, \quad (14)$$

$$\mathbf{h}_2(\boldsymbol{\theta}(t)) = \boldsymbol{\theta}(t) - \boldsymbol{\theta}_{\min} \geq 0, \quad (15)$$

$$\mathbf{h}_3(\boldsymbol{\theta}(t)) = \boldsymbol{\theta}(t) - \boldsymbol{\theta}_{\max} \leq 0. \quad (16)$$

Here \mathbf{f} is identified kinematic model, $\mathbf{X}_{i,\text{des}}$ is the i th desired pose element for the inverse kinematics, n_{dof} is the degrees of freedom that are used for the end-effector control, \hat{n}_{eff} are the end-effector pose elements with a soft constraint, n_{eff} are end-effector pose elements with hard constraints. We assume $\epsilon = 10^{-3}$.

The following are free parameters which the motion designer can tune in order to bias the motion.

- (1) The reference angles for the joints, $\theta_i^{\text{ref}}(t)$. We choose $\theta_i^{\text{ref}}(t) = \theta_i^{\text{ref}}$ (constant values), which is the joint angles corresponding to the robot pose shown in Fig. 1.
- (2) The joint weights, w_i . We intuitively chose a weight of 1 for the joints for the hand degrees of freedom and choose a weight of 10 for the joints in the mid- and lower-body. This choice of this particular weight distribution has the effect of finding solutions that involve bigger excursions of the hand degrees of freedom than the body degrees of freedom, similar to what humans would do when doing

tasks using their hands. Alternately, the weights can be chosen using inverse optimization using motion capture data (e.g., See Liu et al. ²⁸).

- (3) The end-effector weights, W_j . These weights are only used in the tray task for the orientation of left hand. They were tuned to 10.

5.2. Iterative Learning Control (ILC) Algorithm

The joint angle solution $\theta(t)$ leads to poor performance when we implement it on the robot because of the imperfect kinematic model. So, we implement an iterative learning control to improve the tracking performance of the inverse kinematics solution. We describe the algorithm next.

Let i represent the trial number and j the time index that goes from 1 to n_j (end time). Let the reference motion in task-space be defined by $\mathbf{X}_{\text{ref}}(j)$. Here we have concatenated the position and the orientation in the vector \mathbf{X} . The input to the inverse kinematics solver are the desired poses in end-effector space, which we denote by $\mathbf{X}_{\text{des}}^i(j)$. Let the measured position in task-space be defined by $\mathbf{X}_{\text{act}}^i(j)$. Let the error between actual end-effector position and reference position be denoted by $\mathbf{e}^i(j) = \mathbf{X}_{\text{ref}}^i(j) - \mathbf{X}_{\text{act}}^i(j)$. Our ILC algorithm is shown in Fig. 3 and described below:

- (1) Set the error $\mathbf{e}^0(j) = 0$ and initialize the desired position in the task-space $\mathbf{X}_{\text{des}}^0(j) = \mathbf{X}_{\text{ref}}(j)$.
- (2) For subsequent trials do:
 - Command modification in end-effector space: Update the feed-forward position command using the tracking error at trial i , $\mathbf{X}_{\text{des}}^i(j) = \mathbf{X}_{\text{des}}^{i-1}(j) + \mathbf{\Gamma}\mathbf{e}^{i-1}(j)$. The manually tuned learning gain, $\mathbf{\Gamma}$, is a 6x6 matrix of the form; $\mathbf{\Gamma} = \text{diag}\{\gamma_1, \gamma_2, \dots, \gamma_6\}$. Further, for ILC to converge we need $0 < \gamma_i \leq 1$.
 - Command initialization in joint-space: For the desired position in task-space $\mathbf{X}_{\text{des}}^i(j)$, use the inverse kinematics solver to find a desired joint command $\theta^i(j)$.
 - Command execution on robot: Send the feed-forward commands $\theta^i(j)$ ($j = 1, 2, \dots, n_j$) to the low level controller. Save the resulting tracking errors in the end-effector space, $\mathbf{e}^i(j)$ ($j = 1, 2, \dots, n_j$).
- (3) Stop when the error metric e_{norm}^i does not improve between trials. The learnt feed-forward command is then $\theta^i(j)$ ($j = 1, 2, \dots, n_j$).

The error metric to check convergence is given as follows

$$e_{\text{norm}}^i = \frac{1}{n_j} \sum_{j=1}^{n_j} \sum_{k=1}^{n_{\text{eff}}} (e_k^i(j))^2, \quad (17)$$

where $e_k^i(j)$ is the tracking error in the pose element k , at iteration i and at time j , $n_{\text{eff}} = 6$ and n_j is the total data points in the trial.

Tuning the learning gain: The free parameters in the ILC algorithm are the learning gains, $\mathbf{\Gamma}$. The choice of this parameter affects the learning rate, convergence, and stability (see Bhounsule et al.¹² for conditions on the learning gain). We simplify the choice, by using the same learning parameter for all six degrees of freedom. Thus $\mathbf{\Gamma} = \gamma I$, where I is 6x6 identity matrix and $0 < \gamma \leq 1$.

Zero phase filtering: We use a zero phase filter to remove sensor noise and to provide a zero phase lag²⁹. The zero phase filtering is done as follows. We first filter in the forward direction using a second order Butterworth filter with cutoff frequency of 1 Hz. Next, we pad the forward filtered signal with about 120 reflected data points at the beginning and at the end. Then, we reverse the concatenated signal and filter again with the same Butterworth filter. This process of forward filtering followed by reverse filtering produces a signal with zero delay. The padding of data removes unnecessary transients in the beginning and the end of the filtered signal. Note that the zero-phase filtering is anti-causal and it needs the sensor values for the entire trajectory and is done offline.

5.3. Results for task-level iterative learning control

We demonstrate the implementation of our algorithm on the humanoid robot on three control tasks (see Fig. 4): (1) drawing the figure eight (writing task), (2) using one hand to serve a glass of drink (glass task), and (3) serving a drink atop a tray using both hands (tray task). For the writing and glass task, the joints 1 through 7 in the body, and 19 through 26 in the right hand were used. For the tray tasks, all joints except the joints controlling the head (joints 8, 9, and 10) were used. Note that the joints A through E are neither sensed or actuated but need to move in order to allow the loop joint in the lower body to move. Also see Fig. 1 (b) and (c) for the experimental set up and placement of markers for learning.

Task 1: Writing Task: The writing task consists drawing the lemniscate, which is the figure eight or the ∞ symbol. We specify the lemniscate equation in $\mathbf{X}_{\text{des}}(t)$. We use the inverse kinematic solver (see Sec 5.1) to compute the joint position command $\boldsymbol{\theta}(t)$. Here $n_{\text{eff}} = 6$ corresponding to the position and orientation of the right hand and $\hat{n}_{\text{eff}} = 0$. The joint command is then played back on the robot. The motion of the end-effector is tracked by three markers placed on the right hand. The kinematic model does not take into account the actuator dynamics or the link deflection, and does not produce error free tracking. Finally, we use the the inverse ILC algorithm to improve the tracking performance.

The learning parameter γ was manually tuned to 0.3 by running a few trials on the robot. While the robot is learning, the robots draws in the air and the motion capture system helps to measure the motion of the end-effector. The ILC algorithm uses the tracking error to improve the performance. The error metric (see Eqn. 17) for Trial 1 is 3×10^{-2} , and it decreases to 2×10^{-5} in 18 trials. This is almost a

three orders of magnitude improvement. The rate of convergence is shown in Fig. 5 (a).

Fig. 6(a) shows errors in the position and the orientation as a function of time for the first trial and the converged trial. We can see that the error is reduced to almost zero by the learning algorithm. Fig. 7 (a) and (b) shows a plot of the drawing task in the x-y (horizontal plane) and the x-z plane (fore-aft plane). The solid black line is the reference. The dash-dotted line is the robot motion during trial 1. The dashed lined shows the converged motion at trial 18.

After the robot has learnt the motion, we made the robot draw the lemniscate on a piece of paper using a brush dipped in black paint. A snapshot of the robot after completing the task is shown in Fig. 4 (a) and the actual drawing is shown inset.

Task 2: Glass Task The glass task consists of moving glass in a straight line in the fore-aft direction while maintaining a constant height and constant orientation throughout the motion, as if the robot is serving a drink to a person in front of it. Fig. 4 (b) shows the task executed with the converged trial and with the glass filled with a liquid.

We use the same value for the learning parameter, $\gamma = 0.3$. Also, $n_{\text{eff}} = 6$ corresponding to the position and orientation of the right hand and $\hat{n}_{\text{eff}} = 0$. We did learning in three scenarios and we describe them next. For all the learning scenarios, the motion capture system tracked a set of three markers placed on top of the glass. The three markers were used by the motion capture software to give the glass position and orientation for the learning algorithm.

Learning from identified model: We initialize the learning from the kinematics model we identified. Trial 1 produced an error of 2×10^{-2} (see Eqn. 17). This is understandable because our forward model ignores the actuator dynamics and the link deflection. However, after using iterative learning control the error norm decreased to 2×10^{-4} in 8 trials. This is about a two orders of magnitude improvement. A plot of the convergence rate is shown as filled circles in Fig. 5 (b).

Learning incorrect potentiometer calibration: In order to test if our learning algorithm can learn from incorrect potentiometer calibration, we change the gain and the bias by 10% on joint 7 (bends the torso in the fore-aft direction), joint 20 (moves the right shoulder in the front back direction), and joint 23 (right elbow joint) (see Fig. 1). We use the earlier kinematic model that does not account for the incorrect potentiometer calibration. The error in Trial 1 was 5×10^{-2} but it decreased to 7×10^{-4} in 9 trials. This is almost two orders of magnitude improvement. A plot of the convergence rate is shown with diamond shape in Fig. 5(b). Also, Fig. 6 (b) shows the error in tracking for Trial 1 and for Trial 9.

Learning when some joints are unresponsive: One situation that can come when such robots are deployed is that one or more joints might be unresponsive in the event of a controller failure. In such cases, it is desirable to be able to temporarily accommodate the failure. The ILC framework can be used for temporary fault accommodation by exploiting the system kinematic redundancy.

To simulate a fault, we made the controller on two joints unresponsive. The two joints we made unresponsive were the joint Pelvis joint (bends the mid body in the fore-aft direction) and Torso Forebend (bends the upper body in the fore-aft direction) (see Fig. 1). The system first runs the converged command from the first case presented in this section (Trial 1 shown by the cross in Fig. 5(b)). This gives a rather large error (about 4×10^{-2}) in the task space. We set new joint limits in the inverse kinematics solver and use the ILC algorithm to improve the tracking errors in task space. In 9 trials, the ILC algorithm converges and produces the net error of 6×10^{-5} . This is almost two orders of magnitude improvement. The crosses in Fig. 5(b) shows the convergence rate.

Fig. 8 shows the joint command sent to the robot for all the 14 joints for this experiment. The dashed lines is the initial command that assumes all joints are working fine. The solid lines is the learnt command. Note that the Elbow joint mostly compensates for the unresponsive Pelvis and Torso Forebend joints.

Task 3: Tray Task The tray task consists of the robot moving a tray with both hands in a straight line in the fore-aft direction while maintaining a constant height and constant orientation of the *tray* throughout the motion. Fig. 4 (c) shows the task executed with the learnt command and with a filled glass placed on the tray.

There were practical issues in getting the tray task to work. The robot is designed such that the robot hands can get within about 75 cm from each other. Thus the robot needs a suitably wide tray. The ideal way to move an object with both hands is to maintain the position of both hands with respect to each other and a constant orientation of the hands. This ensures that the tray moves level with respect to the world frame. However, the robot joints limits are quite severe that we could not find inverse kinematics solution that will give a big range of motion for the tray while respecting position and orientation for both hands. We modified the tray supports in such a way that the robot makes line contact with the tray on the left hand and point contact with the tray on the right hands (see 4 (c)). This way we do not need to do orientation control on the right hand. This eases the hard constraint on the right hand letting us do a large range of motion.

For this task, we are interested in maintaining the orientation and position of the left hand and the position of the right hand. Thus $n_{\text{eff}} = 9$ which are the hard constraints in Eqn. 14. We try to maintain a suitable orientation of the right hand using soft constraint, thus $\hat{n}_{\text{eff}} = 3$. We tuned the weight $W_j = 10$ in Eqn. 13. The rest pose given by θ_i^{rest} is shown in Fig. 1 (c). The learning parameter γ was set at 0.3 for all end-effector elements. The positions were measured in meters and the orientations were measured in terms of quaternions. While the robot is learning, the

tray is not placed on the hands. The tracking errors are used by the ILC algorithm to improve the performance.

The error metric (see Eqn. 17) for Trial 1 was 2×10^{-2} and 2×10^{-3} for the left hand and right hand respectively. This decreased to 3×10^{-4} and 3×10^{-5} for the left hand and right respectively in 14 trials. This is almost three orders of magnitude improvement for each hand individually. The rate of convergence is shown in Fig. 5 (c).

Fig. 6(c) shows errors in the position and orientation as a function of time for the first trial and the converged trial for the right and left hand. We can see that the error is reduced to almost zero by the learning algorithm. Note that the orientation error for the right is non-zero because we did not learn this error. We enforced the right hand orientation as a soft constraint in the inverse kinematics.

After the robot has learnt the motion, we checked if the ILC algorithm indeed works by placing a tray with a glass filled with liquid. A snapshot of the robot after completing the task is shown in Fig. 4 (c).

6. Discussion

We have provided a comprehensive technique for accurate task space tracking for a humanoid robot with serial and parallel linkage using a motion capture (MoCap) system. We started with an approximate kinematic model and identified its parameters using non-linear least squares using a constraint optimization software. The average error in the position and orientation was 1 cm and 1.84 degrees respectively. The limited accuracy was because of joint flexibility and sensor noise which was not factored in our rigid body kinematics model. Thus, using the kinematics model, we were unable to get accurate task space tracking.

We then improved on the tracking by incorporating an iterative learning control algorithm that corrected the joint command based on errors in end-effector tracking. We tested the algorithms on a drawing task, a serving a drink task, both involving one hand and serving a drink using a tray using both hands. In each of these cases, the task space error was reduced by at least two to three orders of magnitude in a maximum of 18 trials.

Practical kinematic calibration for humanoid robots: Humanoid robots, unlike traditional industrial manipulators, might need periodic calibration for the following reasons: (1) Some humanoids (e.g., the robot used in this research) are retrofitted with replacement parts during their life cycle; (2) To keep cost low, humanoids are often fitted with cheap sensors (e.g., potentiometers) which need frequent calibration; (3) Some humanoids have special transmissions such as pulleys, cables, and tendons which need to be frequently calibrated; and (4) Sometimes CAD models of humanoids may not be available. Hence a quick and accurate technique for calibration is needed.

The marker-based motion capture system allows us to collect rich data from

multiple end-effectors simultaneously. In our study, the motion capture time lasted for 30 minutes for a 26 degree of freedom humanoid. Starting from a crude model of the robot which was estimated using a protractor and a ruler, our non-linear least squares optimization program was able to fit 178 parameters in about 1 hour on a standard desktop computer (circa 2010). Our fit accuracy was 1 cm for the position, which is within 0.6% when compared with the robot height, and 1.84 degrees for the orientation, averaged across all the end-effectors. Thus kinematic calibration procedure we propose is fast, accurate and can be automated if desired.

Better models give better convergence: Though ILC scheme is designed to correct for modeling errors, better models give better convergence³⁰. This is clear when we compare the two test cases in the glass task; (1) learnt from an identified model vs. (2) learnt from identified model but with incorrect pot calibration. In the former case, the error in the converged trial is 2×10^{-4} which is smaller than the error in the converged trial of 7×10^{-4} in the latter case.

Another place where better models lead to better convergence is evidenced when we compare the writing task with the glass task. The error in the converged trial for the writing task and glass task are 2×10^{-5} and 2×10^{-4} respectively. Thus in the writing task, the final convergence is an order of magnitude better than in the glass task. This is explained as follows: In the glass task, the robot has to stretch its hands and body outward to reach out. The robot is flexible and the structural loading causes link deflection which is not accounted in our model. On the other hand, in the writing task, the robot does not have to reach out and the link deflection is much smaller. In other words, our kinematic model for the writing task is much more accurate than in the glass task. The evidence for the above explanation can be seen by comparing the error in the z position between the writing task shown in Fig. 6 (a) (iii) and the glass task shown in Fig. 6 (b) (iii). It is seen that the error in the z-direction keep increasing as the robot stretches its hand to move the glass to serve the drink.

Use of existing robot model: Our method treats the inverse kinematics solver as a black-box during the learning process. This is advantageous for two reasons: (1) There is no need to rewrite the inverse kinematics solver. This is specially advantageous for humanoid robots that have an inverse kinematics solver already available. (2) Even if the robot model changes a bit, for example due to wear and tear or part replacement, there is no need to re-calibrate the model.

Handling Joint Limits: From our experience, we know that humanoid robots often operate close to the position limits. In our method, the joint limits are handled by the constrained optimization at the inverse kinematics solver (see Fig. 3). In all the experiments reported here, we had multiple joints at their position limits, but the learning proceeded seamlessly, converging to a small tracking error. We believe that this is a significant advantage of our method over others that work at the

velocity level and use the Jacobian to map from task-space to joint space.

Joint limits: Another way of doing the learning is to do the update in the joint space^{10,11}

$$\theta_{\text{des}}^i(j) = \theta_{\text{des}}^{i-1}(j) + \gamma \hat{J}^{-1} e^{i-1}(j), \quad (18)$$

The disadvantage of the above update rule is that if one is close to the joint limit than the projection via the Jacobian \hat{J}^{-1} may not give an improvement at all. This might cause the ILC algorithm to not converge, or even worst, to become unstable. Please see Bhounsule et al.¹² showing comparizon of our method with that of Arimoto^{10,11}. Our method does not suffer from this because we use a nonlinear inverse that find solutions within the joint limits.

Fault accommodation: Fault diagnosis and isolation is critical to reduce downtime and also for timely intervention. Once the fault is detected, the ILC algorithm we present here can be used to accommodate failures by exploiting the robot kinematic redundancy. We have presented one such scenario in this paper: the robot serves the glass without spilling even though two of its joints are unresponsive. A future direction would be to exploit the structural redundancy (redundant hardware, e.g., two motors on one joint) and the functional redundancy (hardware that performs different functions, e.g., estimate position using inertial measurement unit if the potentiometer breaks)³¹ to develop work-arounds when joints on the robot become unresponsive.

Limitations of our method: Our method has all the limitations of ILC: it is an offline method; it needs manual tuning to work well; and it can only improve a single trajectory at a time. In addition, our method needs an inverse kinematics solver that is able to find solutions within joint limits. This can be computationally expensive and can lead to issues if the manipulator is in singular configuration.

Appendix A. Multimedia Extension

A video of the robot doing the three tasks is available as a multimedia extension at the following url:

<https://www.disneyresearch.com/publication/accurate-task-space-tracking/>

References

1. K. J. Astrom, Pid controllers: theory, design and tuning, *Instrument society of America*, 1995.
2. T. B. Cunha, C. Semini, E. Guglielmino, V. J. De Negri, Y. Yang, and D. G. Caldwell, Gain scheduling control for the hydraulic actuation of the hyq robot leg, in *Proc. of COBEM*, 2009.
3. D. Sun and J. K. Mills, High-accuracy trajectory tracking of industrial robot manipulator using adaptive-learning scheme, in *American Control Conference*, vol. 3, pp. 1935–1939, IEEE, 1999.

4. A. Guez and Z. Ahmad, Solution to the inverse kinematics problem in robotics by neural networks, in *IEEE International Conference on Neural Networks*, pp. 617–624, IEEE, 1988.
5. R. Köker, C. Öz, T. Çakar, and H. Ekiz, A study of neural network based inverse kinematics solution for a three-joint robot, *Robotics and Autonomous Systems*, vol. 49, no. 3, pp. 227–234, 2004.
6. M. I. Jordan and D. E. Rumelhart, Forward models: Supervised learning with a distal teacher, *Cognitive science*, vol. 16, no. 3, pp. 307–354, 1992.
7. E. Oyama, N. Y. Chong, A. Agah, and T. Maeda, Inverse kinematics learning by modular architecture neural networks with performance prediction networks, in *IEEE International Conference on Robotics and Automation*, vol. 1, pp. 1006–1012, IEEE, 2001.
8. S. Arimoto, S. Kawamura, and F. Miyazaki, Bettering operation of robots by learning, *Journal of Robotic systems*, vol. 1, no. 2, pp. 123–140, 1984.
9. D. A. Bristow, M. Tharayil, and A. G. Alleyne, A survey of iterative learning control, *IEEE Control Systems*, vol. 26, no. 3, pp. 96–114, 2006.
10. S. Arimoto, M. Sekimoto, and S. Kawamura, Iterative learning of specified motions in task-space for redundant multi-joint hand-arm robots, in *IEEE International Conference on Robotics and Automation*, pp. 2867–2873, IEEE, 2007.
11. S. Arimoto, M. Sekimoto, and S. Kawamura, Task-space iterative learning for redundant robotic systems: existence of a task-space control and convergence of learning, *SICE Journal of Control, Measurement, and System Integration*, vol. 1, no. 4, pp. 312–319, 2008.
12. P. Bhounsule, K. Yamane, and A. Bapat, A task-level iterative learning control algorithm for accurate tracking in manipulators with modeling errors and stringent joint position limits, in *ASME Dynamics Systems and Controls Conference, Minneapolis, MN, USA*, 2016.
13. P. Bhounsule and K. Yamane, Iterative learning control for accurate task-space tracking with humanoid robots, in *International Conference on Humanoid Robots, Seoul, South Korea*, 2015.
14. Natural Point, Optitrack - arena - body motion capture software for 3d character animation and more. World Wide Web electronic publication, 2013.
15. B. Mooring, The effect of joint axis misalignment on robot positioning accuracy, in *Proceedings of the ASME international computers in engineering conference*, pp. 151–156, 1983.
16. B. Mooring and G. Tang, An improved method for identifying the kinematic parameters in a six axis robot, in *Proceedings of the ASME International Conference on Computers in Engineering*, 1984.
17. K. Gupta, Kinematic analysis of manipulators using the zero reference position description, *The International Journal of Robotics Research*, vol. 5, no. 2, pp. 5–13, 1986.
18. S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, C. T. John, E. Guendelman, and D. G. Thelen, Opensim: open-source software to create and analyze dynamic simulations of movement, *IEEE Transactions on Biomedical Engineering*, vol. 54, no. 11, pp. 1940–1950, 2007.
19. B. W. Mooring, Z. S. Roth, and M. R. Driels, *Fundamentals of manipulator calibration*. Wiley-interscience, 1991.
20. J. Denavit, A kinematic notation for lower-pair mechanisms based on matrices., *Trans. of the ASME. Journal of Applied Mechanics*, vol. 22, pp. 215–221, 1955.
21. A. Broun, C. Beck, T. Pipe, M. Mirmehdi, and C. Melhuish, Building a kinematic

- model of a robots arm with a depth camera, in *Advances in Autonomous Robotics*, pp. 105–116, Springer, 2012.
22. M. G. Hollars, D. E. Rosenthal, and M. A. Sherman, Sd/fast users manual, *Symbolic Dynamics Inc*, 1991.
 23. P. E. Gill, W. Murray, and M. A. Saunders, Snopt: An sqp algorithm for large-scale constrained optimization, *SIAM journal on optimization*, vol. 12, no. 4, pp. 979–1006, 2002.
 24. R. P. Judd and A. B. Knasinski, A technique to calibrate industrial robots with experimental verification, *Robotics and Automation, IEEE Transactions on*, vol. 6, no. 1, pp. 20–30, 1990.
 25. J.-M. Renders, E. Rossignol, M. Becquet, and R. Hanus, Kinematic calibration and geometrical parameter identification for robots, *IEEE Transactions on Robotics and Automation*, vol. 7, no. 6, pp. 721–732, 1991.
 26. H. A. Park, M. A. Ali, and C. G. Lee, Closed-form inverse kinematic position solution for humanoid robots, *International Journal of Humanoid Robotics*, vol. 9, no. 03, p. 1250022, 2012.
 27. L. C. Wang and C.-C. Chen, A combined optimization method for solving the inverse kinematics problems of mechanical manipulators, *IEEE Transactions on Robotics and Automation*, vol. 7, no. 4, pp. 489–499, 1991.
 28. C. K. Liu, A. Hertzmann, and Z. Popović, Learning physics-based motion style with nonlinear inverse optimization, *ACM Transactions on Graphics (TOG)*, vol. 24, no. 3, pp. 1071–1081, 2005.
 29. H. Elci, R. W. Longman, M. Q. Phan, J.-N. Juang, and R. Ugoletti, Simple learning control made practical by zero-phase filtering: Applications to robotics, *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 49, no. 6, pp. 753–767, 2002.
 30. C. Atkeson and J. McIntyre, Robot trajectory learning through practice, in *IEEE International Conference on Robotics and Automation*, vol. 3, pp. 1737–1742, IEEE, 1986.
 31. M. L. Visinsky, J. R. Cavallaro, and I. D. Walker, Robotic fault detection and fault tolerance: A survey, *Reliability Engineering & System Safety*, vol. 46, no. 2, pp. 139–158, 1994.

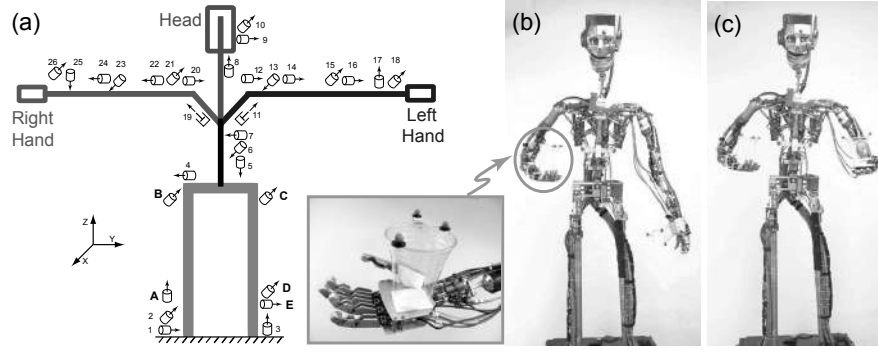


Fig. 1. (a) The kinematic model of the humanoid robot in the zero reference pose has 26 degrees of freedom. In the reference pose all the joint angles zero. The names of these joints are; 1) **Body Forebend (BF)**, 2) **Body Sway (BS)**, 3) **Body Twist (BT)**, 4) **Pelvis (P)**, 5) **Torso Twist (TT)**, 6) **Torso Sway (TS)**, 7) **Torso Forebend (TF)**, 8) **Head Turn (HT)**, 9) **Head Nod (HN)**, 10) **Head Tilt (HTLT)**, 11) **Left Shoulder Shrug (LSS)**, 12) **Left Arm Forebend (LAF)**, 13) **Left Arm Out (LAO)**, 14) **Left Arm Swing (LAS)**, 15) **Left Elbow (LE)**, 16) **Left Wrist Twist (LWT)**, 17) **Left Wrist In Out (LWIO)**, 18) **Left Wrist Forward Back (LWFB)**, 19) **Right Shoulder Shrug (RSS)**, 20) **Right Arm Forebend (RAF)**, 21) **Right Arm Out (RAO)**, 22) **Right Arm Swing (RAS)**, 23) **Right Elbow (RE)**, 24) **Right Wrist Twist (RWT)**, 25) **Right Wrist In Out (RWIO)**, 26) **Right Wrist Forward Back (RWFB)**. The joints named **A**, **B**, **C**, **D**, **E** are neither actuated nor sensed. Gross robot specifications are: height = 1.8 m (5'11"), body width = 0.28 m (9"), and length of hand = 0.67 m (2'2"). (b) A photo of the humanoid robot that we used in this study. The above pose is our reference pose for the inverse kinematics (see Sec 5.1) for the writing and glass task. A plastic glass is glued to the right hand of the robot and is shown inset. We put three markers on the top of the glass which we track using our motion capture system. (c) The rest pose for the tray task. There are three markers on top of glasses attached to the hands. We use the glasses to elevate the markers so that they are in clear view of the motion capture system.

Table 1. **Fit accuracy:** Root mean square error for the position and the orientation for the head, the left hand and the right hand. See Section 4.4 for details.

Name	Head	Left Hand	Right Hand	Average Root Mean Square Error
Δx in cm	0.47	0.71	0.74	0.64
Δy in cm	0.28	0.42	0.56	0.42
Δz in cm	0.57	0.76	0.64	0.66
Δr in cm	0.80	1.12	1.12	1.00
$\Delta\theta_x$ in degrees	0.67	0.97	1.15	0.93
$\Delta\theta_y$ in degrees	0.8	1.22	1.56	1.19
$\Delta\theta_z$ in degrees	1.03	1.07	0.97	1.02
$\Delta\theta$ in degrees	1.47	1.89	2.17	1.84




Fig. 2. **Distribution for the fit accuracy for the left hand:** Histogram of the position error are on the left-side and the orientation error are on the right-side. There are about 750 data points per histogram.

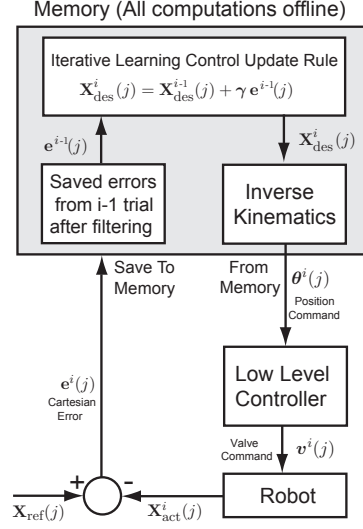


Fig. 3. Block diagram of our end-effector space iterative learning control algorithm. We filter the cartesian space errors $e^{i-1}(j)$ using a zero phase filter before applying the iterative learning update rule. The super-script i denotes the trial number and j denotes a time instance.




Fig. 4. (a) Robot doing the writing task. The figure "eight" drawn by the robot is shown inset. (b) Robot doing the glass task. The robot serving the drink without spilling. The learning was done with no liquid in the glass (c) The robot doing the tray task with a drink placed on the tray. The learning was done without the tray.

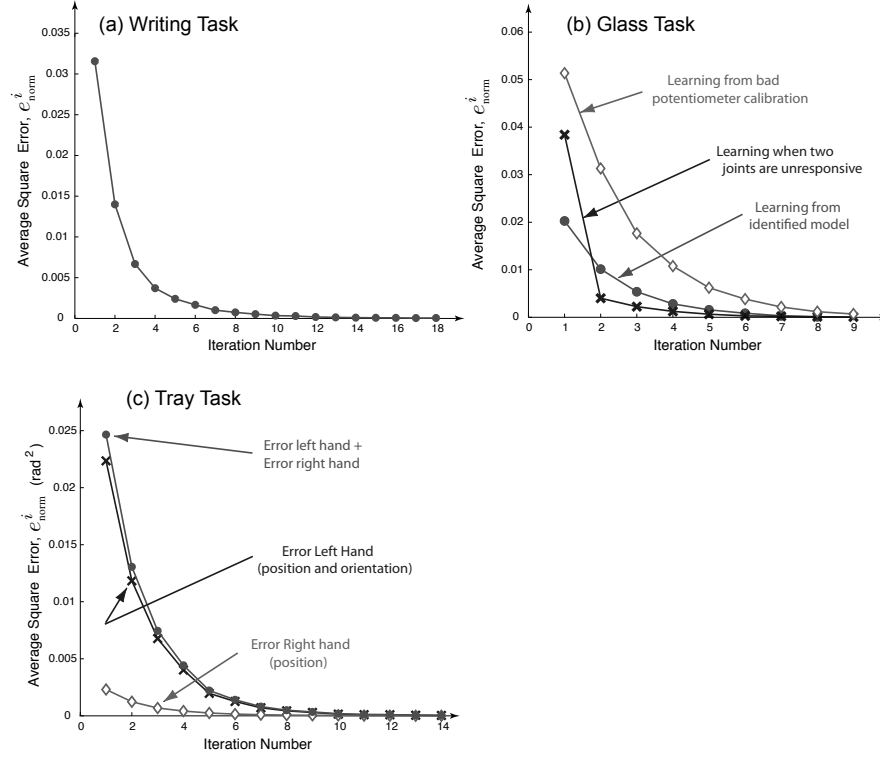


Fig. 5. Convergence of errors as a function of iteration number for (a) the writing task, and (b) the glass task, (c) the tray task

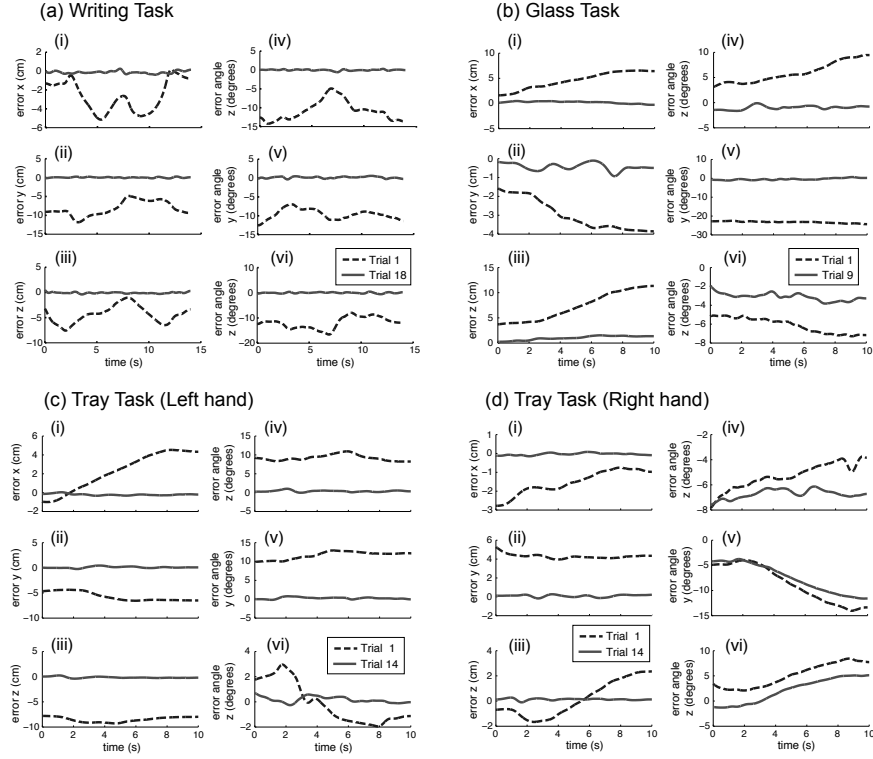


Fig. 6. Position and Orientation Errors for (a) the writing task and (b) the glass task (learning incorrect potentiometer calibration) (c) Tray task, left hand, and (d) Tray task, right hand. (i,ii,iii) Errors in the position in the x, y, z directions in cms and (iv,v,vi) Errors in the angle in degrees in the x, y, z direction. x is in the fore-aft direction, y is in the robot left-right direction and z is in the up-down direction. We do not learn orientation error in the left hand. This is enforced as a soft constraint in the inverse kinematics.

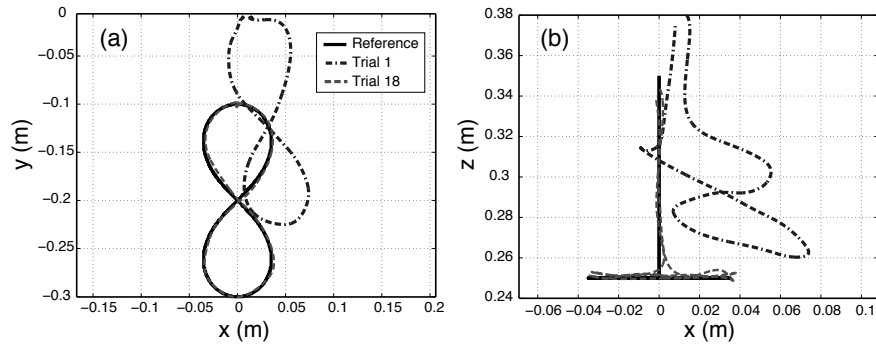


Fig. 7. (a) Plot in the x-y plane (z axis lines up with gravity). (b) Plot in the x-z plane (fore-aft plane of robot). The dash-dot lines and dashed lines are the marker position measured by the motion capture system.

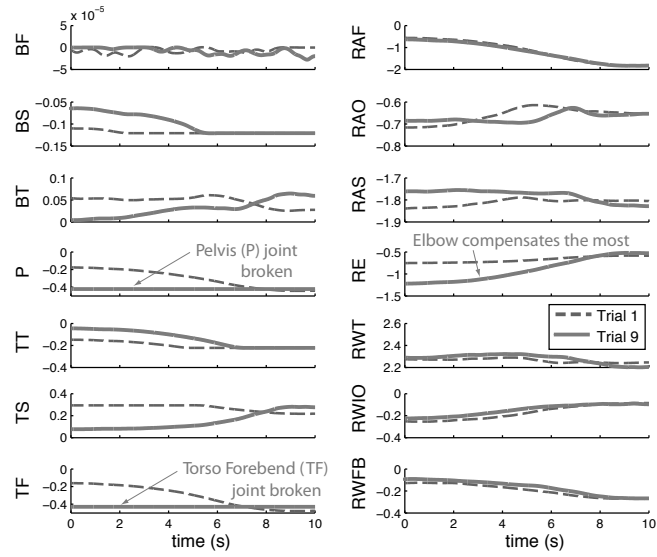


Fig. 8. Joint Command vs. Time for normal operations (dashed lines) and for learnt motion after accommodating for unresponsive joints (continuous lines). The Pelvis Joint (P) and the Torso Forebend (TF) both are frozen and will not move. The Right Elbow (RE) joint compensates the most for task space control.

Biography



Fig. 9. Pranav A. Bhounsule

Dr. Pranav A. Bhounsule is an Assistant Professor in the Department of Mechanical Engineering at the University of Texas at San Antonio. He received his BE degree from Goa University in 2004, the MTech degree from Indian Institute of Technology Madras in 2006, and the Ph.D. degree from Cornell University in 2012. From 2012 to 2014, he was a post doctoral research at Disney Research Pittsburgh. He received the award for Best paper in Biologically Inspired Robotics at the Climbing and Walking Robots Conference 2012. He also received a US National Science Foundation Research Initiation grant in 2016. His research interest includes model-based and learning-based control methods for legged and humanoid robots.



Fig. 10. Katsu Yamane

Dr. Katsu Yamane is a Senior Research Scientist at Disney Research, Pittsburgh and an Adjunct Associate Professor at the Robotics Institute, Carnegie Mellon University. He received his B.S., M.S., and Ph.D. degrees in Mechanical Engineering in 1997, 1999, and 2002 respectively from the University of Tokyo, Japan. Prior to joining Disney in 2008, he was an Associate Professor at the University of Tokyo and a postdoctoral fellow at Carnegie Mellon University. Dr. Yamane is a recipient of King-Sun Fu Best Transactions Paper Award and Early Academic Career Award from IEEE Robotics and Automation Society, and Young Scientist Award from Ministry of Education, Japan. His research interests include humanoid robot control and motion synthesis, physical human-robot interaction, character animation, and human motion simulation.