# Exploiting Vestibular Output during Learning Results in Naturally Curved Reaching Trajectories

**Ganghua Sun and Brian Scassellati**
Computer Science Department, Yale University
New Haven, Connecticut 06511, United States
{ganghua.sun, brian.scassellati}@yale.edu

## Abstract

Teaching a humanoid robot to reach for a visual target is a complex problem in part because of the high dimensionality of the control space. In this paper, we demonstrate a biologically plausible simplification of the reaching process that replaces the degrees of freedom in the neck of the robot with sensory readings from a vestibular system. We show that this simplification introduces errors that are easily overcome by a standard learning algorithm. Furthermore, the errors that are necessarily introduced by this simplification result in reaching trajectories that are curved in the same way as human reaching trajectories.

## 1. Introduction

Redundancy in the kinematic structure of the human body enables the same sensorimotor task to be executed in different ways. However, this flexibility comes at a cost. Bernstein first recognized that the large number of degrees of freedom (DOFs) in the kinematic structure make the learning of sensorimotor tasks a very hard problem (Bernstein, 1967). In statistical learning theory, this same problem has been called 'the curse of dimensionality'.

Reaching is one of the most thoroughly investigated sensorimotor tasks. Because the human arm contains more joints than there are dimensions in the work space, there are in principle infinite postures that reach the same target in space. The experiments carried out by Morasso showed the remarkable result that human reaching movements invariably possess two characteristics: the trajectory is gently curved and the velocity profile is bell-shaped (Morasso, 1981). Many theories have been proposed to explain the observations made by Morasso with the assumption that motor learning is governed by some optimization measure. The optimization measures suggested include minimum jerk (Hogan, 1984), maximum smoothness (Flash and Hogan, 1985), minimum torque (Uno et al., 1989) and, recently, minimum variance of the final hand position (Harris and Wolpert, 1998). For a comprehensive review, please refer to (Todorov, 2004).

This paper offers a new hypothesis to explain the curvature in reaching trajectories from a different perspective. Guenther and Barraca observed that the curvature in reaching trajectories can be a side effect of motor learning due to a bias in the training data (Guenther and Barreca, 1997). This bias leads to the desirable consequence of singularity avoidance. We hold the same view that trajectory curvature can be just a side effect of motor learning. But we believe that the cause of this side effect is more fundamental: it is the simplification made by the learning system to reduce the dimensionality of the reaching problem that leads indirectly to curved trajectories.

We have shown previously that a humanoid robot can learn the task of reaching with a very modest number of training samples (Sun and Scassellati, 2004). One major weakness of our approach was that the head of the robot needed to remain stationary during learning. If the head posture were changed, a new round of training would be required. In this paper, we extend our previous approach so that all DOFs in the arm and the neck are involved during learning, but the learned forward model has the same input and output dimensions as before. With the aid of visual feedback during the arm movement, reaching accuracy remains the same as previously achieved. The only apparent change is that reach trajectories that were previously straight (in Cartesian coordinates) are now gently curved. This observation serves as the foundation of our hypothesis to explain the curvature of human reach trajectories.

The next section describes the original system we used for learning to reach and provides arguments for its biological plausibility. In Section 3, we extend the original system to allow simultaneous activation of all motors in the arm and the neck during motor babbling and explain the cause of the resulted gentle curvature in the reaching trajectories. A discussion is given in Section 4. Section 5 concludes this paper by summarizing our contributions.

## 2. The Basic System

Nico is an upper-torso humanoid robot developed to match the body dimension of an average one-year-old infant. It has four DOFs in the neck and six DOFs in each arm. A gyroscope is mounted in the head on top of all neck joints. Fig. 1(A) shows a dimetric view of Nico. Fig. 1(B) shows the kinematics structure of the neck joints and the joints in the right arm viewed from behind.

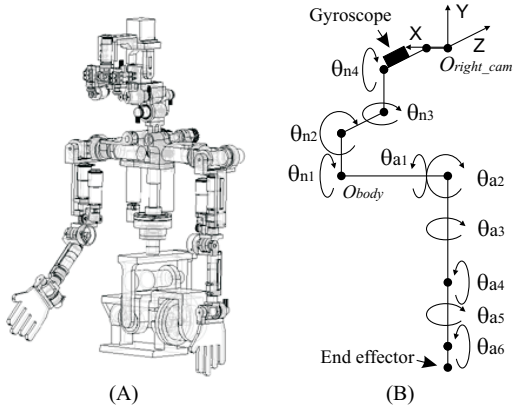


Figure 1: (A) A dimetric view of Nico, an upper-torso humanoid designed to match the size of an average one-year-old infant. (B) The kinematic structure of the neck and the right arm viewed from behind. The origins of the eye and the body-centered coordinate system are indicated with $O_{right\_cam}$ and $O_{body}$ respectively.

In (Sun and Scassellati, 2004), we introduced a learning-based approach that enables Nico to incrementally generate reaching trajectories. The approach can be summarized by the diagram in Fig. 2. The forward model shown in the diagram is learned with a Radial Basis Function Network (RBFN) on training samples gathered during a motor babbling phase. It approximates the forward kinematics function $f_{arm} : \theta_{arm} \rightarrow x$ of the arm. $\theta_{arm}$ is a vector that describes the angles of the arm joints and $x$ is a vector that describes the corresponding 3D position of the end-effector. During a reaching movement, vector $\Delta x$ that describes the moving direction in the task space is computed as $\Delta x = \alpha(x_{target} - x_{ee})$. $x_{ee}$ represents the current position of the end effector. It can be either predicted by the forward model or perceived by the stereo vision system whenever the end effector is visible to both eye cameras. The factor $\alpha$ ensures that the magnitude of $\Delta x$ is equal to a predetermined value of $step\_size$. After $\Delta x$ is computed, it is multiplied by the pseudo-inverse of Jacobian $J$ extracted from the forward model. The result is a vector $\Delta\theta$ that indicates the moving direction in the joint space. The newly generated $\Delta\theta$ is added into $\theta_{arm}$ which is sent to both the arm motors for generating movement and the forward model for updating $x_{pred}$ and $J$. This process is repeated until the end effector of the arm reaches the target.

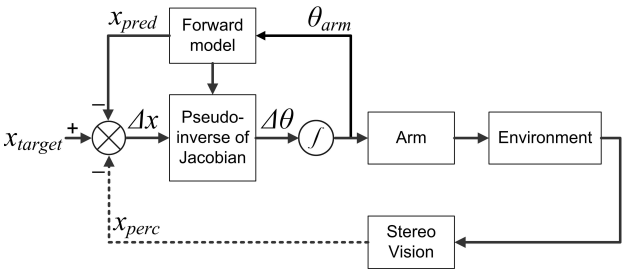


Figure 2: Diagram summarizing our approach of using a learned forward model to incrementally generate reaching trajectory. $x_{pred}$ and $x_{perc}$ stand for the end effector position predicted by the forward model and perceived by the stereo system respectively.

The concept of forward model was first illustrated by (Jordan and Rumelhart, 1992). Jordan and Rumelhart use a learned forward model to train a second network in order to obtain a unique inverse kinematic mapping for a redundant arm. There is growing biological evidence supporting the view that the brain uses forward models for motor control (Wolpert et al., 1995) (Mehta and Schaal, 2002). In contrast to Jordan and Rumelhart's paper, we use the forward model for the purpose of prediction, which is in line with the prevalent opinion about the function of forward models. We adopt a RBFN to represent the forward model partially because it is a universal function approximator that allows for fast training. It has also been suggested as one of the mechanisms used by the brain for motor learning (Pouget and Snyder, 2000). The copy of the vector $\theta_{arm}$ that is sent back to the forward model is called efference copy in the neurophysiological literature. Efference copy is useful in that it avoids most part of the delay of the proprioceptive feedback that has to travel from the periphery to the central nervous system. The proprioceptive feedback from the arm is completely ignored by us as can be seen from Fig. 2 because the motors in our robotic arm always faithfully execute the commands sent to them by the central control. Our approach is closely related to the one proposed by (Bullock et al., 1993). One of the major differences between our approach and theirs is that they learn the direction mapping from $\Delta x$ to $\Delta\theta$ independently while we compute the direction mapping directly from the information in the forward model.

In our original paper, both the neck joints and two wrist joint in the arm were kept fixed during motor babbling. By assigning a large value (110°) to the spread of the Gaussians in the RBFN, only 120 training samples are sufficient to enable a 4-DOF arm to reach accurately. Visual feedback of the end-effector

position can further improve the accuracy. Unpublished follow-up experiments show that when all 6 DOFs in the arm are activated during motor babbling, approximately 400 training samples are needed to learn a forward model of the same quality. Although the increase of the requirement for the training set size is considerable, the total amount of time for gathering these samples is still just around 30 minutes.

## 3. The Extended System

To extend the applicability of our approach to more human-like situations, we now allow the neck joints to move freely during motor babbling. With the neck joints activated, the kinematic function mapping the joint vector into the end-effector position vector becomes: $f_{arm\_neck} : (\theta_{arm}, \theta_{neck}) \rightarrow x$. One way to cope with the expanded kinematic function is simply to increase the number of training samples. However, as has been shown above, the requirement for the number of training samples has more than tripled when we allowed the 2 wrists joints to be activated in addition to the upper 4 DOFs in the arm. If $f_{arm\_neck}$ were to be learned directly, substantially more training samples would be required. Fortunately, this drastic increase of the time to be spent on sample gathering can be avoided by utilizing additional sensory information.

Fig. 3 illustrates a simplified situation where a shift of the head posture leads to changes in both the position and the orientation of the head. The eyes are assumed to stay stationary in relation to the head. The head in its original and shifted posture is painted in black and gray respectively. Although the eye-centered coordinate system's positional change can not be perceived directly, its change in orientation can be sensed by the gyroscope (the vestibular system in case of human). The information delivered by the gyroscope/vestibular system can then be used to correct this orientation shift caused by the posture change of the head, i.e. the coordinate system $OXY$ can be corrected into $OX'Y'$ as shown in Fig. 3.

The original input and output of the kinematics function $f_{arm\_neck}$ are $\theta = (\theta_{arm}, \theta_{neck})$ and

$$x = T^{eye}_{body}(\theta_{neck}) \cdot T^{body}_{ee}(\theta_{arm}) \cdot [0, 0, 0, 1]^t \qquad (1)$$

respectively, where $T^{eye}_{body}$ and $T^{body}_{ee}$ are homogenous transformation matrices. $T^{eye}_{body}$ can be expressed as

$$T^{eye}_{body} = \begin{bmatrix} R_{3\times3} & T_{3\times1} \\ 0_{1\times3} & 1_{1\times1} \end{bmatrix}, \qquad (2)$$

$R$ is the rotation matrix and $T$ is the translation vector. As has been pointed out in the last paragraph, $R$ can be constructed directly from the readings from the gyroscope/vestibular system. We denote $T'^{eye}_{body}$ as

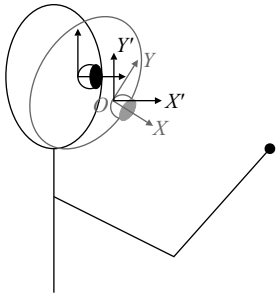


Figure 3: Illustration showing the effect a shifted head posture on the eye-centered coordinate system. The original and the shifted head posture is painted in black and gray respectively. The sensory readings from the gyroscope (the vestibular system in human) allow for the correction of the eye-centered coordination system $OXY$ into $OX'Y'$.

$$T'^{eye}_{body} = \begin{bmatrix} R_{3\times3} & 0_{3\times1} \\ 0_{1\times3} & 1_{1\times1} \end{bmatrix}. \qquad (3)$$

Using $T'^{eye}_{body}$, the perceived position $x$ of the end-effector can be transformed into

$$x' = \left(T'^{eye}_{body}\right)^{-1} \cdot x. \qquad (4)$$

This is similar to correcting $OXY$ into $OX'Y'$ as has been illustrated by Fig. 3.

If we assume that the combined effect of the neck joints is purely rotational so that it is fully eliminated in $x'$, we can learn a forward model with the transformed training samples in the form of $((\theta_{arm})_{input}, (x')_{output})$ instead of $((\theta_{arm}, \theta_{neck})_{input}, (x)_{output})$. In reality, this assumption does not hold true, which can be easily recognized in Fig. 3. Learning a forward model with transformed training samples is equivalent to regarding the translational effect of the neck joints as noise. Through simulations based on the parameters of Nico, we have found out that the standard deviation of the possible positions of the end effector is about three times larger than that of the eye translations caused by the neck joint rotations. Fig. 4 shows a scatter plot for 1000 random positions of the end-effector and a scatter plot for 1000 random positions of the right eye camera. This difference in the ranges of motion between the end-effector and the eye camera means that the position change of the end-effector in the eye-centered coordinate system is caused to a much greater extent by the arm posture than by the head posture. The forward model learned with $((\theta_{arm})_{input}, (x')_{output})$ is not as accurate as the one learned with $((\theta_{arm}, \theta_{neck})_{input}, (x)_{output})$, but it can suffice for the purpose of reaching spatial targets.

Since the body dimensions of Nico and the ranges of motion of its joints closely match those of a one-year-old, the conclusions drawn in this paragraph also apply to a one-year-old.

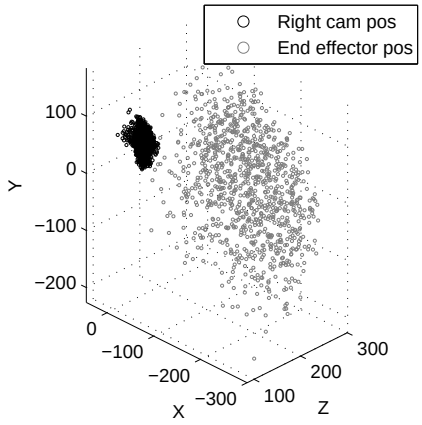


Figure 4: 1000 random positions for both the right eye camera (in black) and the end effector (in gray) are plotted in the body-centered coordinate system. It can be easily seen that the range of motion of the eye camera is much smaller than that of the end effector.

The proposed learning approach dealing with the DOFs in both the head and the arm retains the original dimensionality of problem. It has the additional advantage that the framework shown in Fig. 2 does not need to be changed for generating reaching trajectories. Simulations show that the average positional error of blind reaching is about 20mm. Although this is significantly larger than the average positional error ($<5mm$) of the system that handles only the DOFs in the arm, it can be eliminated by exploiting visual feedback during the reaching movements. When the end effector of the arm is visible to the stereo cameras, the direction vector $\Delta\theta$ can be calculated as

$$\Delta\theta = \alpha J^t (JJ^t)^{-1}(x_{target} - x_{perc}), \qquad (5)$$

If both $x_{target}$ and $x_{perc}$ contain no stereo perception error, the error in $\Delta\theta$ is solely caused by the approximation error in the Jacobian $J$ extracted from the forward model. Since $\Delta\theta$ determines the actual spatial direction the end effector moves along, the error in it causes the actual reaching trajectory deviating from the straight line connecting the starting position of the end effector and the target.

Fig. 5 shows two reaching trajectories produced by simulations. The trajectory that appears to be straight is generated with a forward model $\tilde{f}_{arm}$ that is learned during a motor babbling phase when all 6 DOFs in the arm are activated and the head is kept in a fixed posture. 400 samples are used to train $\tilde{f}_{arm}$. The high quality of this forward model results in a

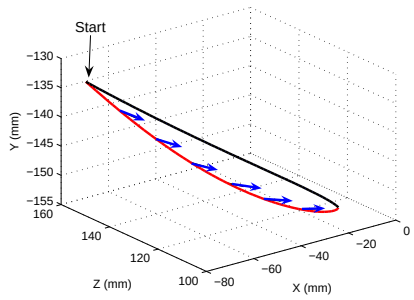


Figure 5: The reaching trajectory that appears to be straight is based on a forward model learned during a motor babbling phase when only the 6 DOFs in the arm are activated. The curved trajectory is based on a forward model learned on samples that use the sensory information of the gyroscope/vestibular system to substitute the proprioception of the neck joints.

reaching trajectory that is apparently straight. The curved trajectory is generated by a forward model $\tilde{f}_{arm\_neck}$ trained on 400 samples gathered during a motor babbling phase when all DOFs in the neck and the arm are activated. These samples take the form of $((\theta_{arm})_{input}, (x')_{output})$ signifying that the sensory feedback of the gyroscope/vestibular system is employed to keep the input dimension of the forward model fixed to 6. Due to the lower quality of $\tilde{f}_{arm\_neck}$ and hence the larger error in $J$ extracted from it, the reaching trajectory based on $\tilde{f}_{arm\_neck}$ is visibly curved. The correct reaching directions on selected points on the trajectory are marked with arrows.

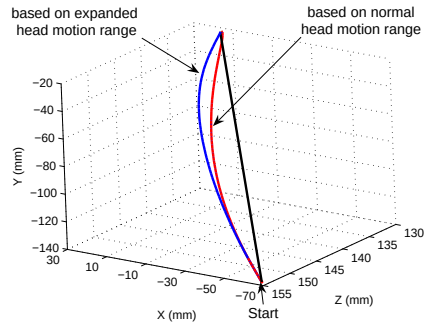


Figure 6: The comparison of these two trajectories show that the curvature of reaching trajecotories is indeed determined by the magnitude of the ignored head translation caused by the neck joint rotations.

To further confirm the conjecture that the curvature of the reaching trajectories based on $\tilde{f}_{arm\_neck}$ is due to the ignored translational effect of the neck joints, we have carried out simulations to compare the trajectories based on two different $\tilde{f}_{arm_neck}$s. For the first model, the neck joints move within their normal ranges of motion when training sam-

ples are being gathered, while for second model, each neck joint moves within a range of motion enlarged by 25%. Such an enlargement increases the average head translation during motor babbling, which in turn leads to larger errors in the second forward model. Fig. 6 displays two trajectories with the same starting end effector and target position. It shows the expected result that the trajectory based on the first forward model is less curved than its counterpart.

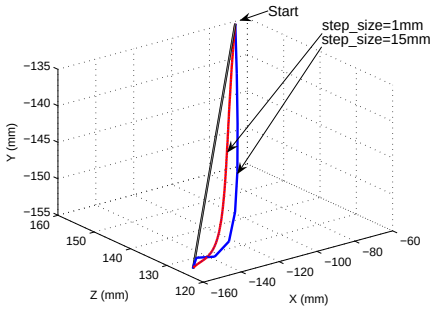


Figure 7: Comparison of two reaching trajectories based on different values of *step_size*. Increasing *step_size* increases the curvature of the trajectory and makes it appear less graceful. However, it also reduces the requirement for the frequency of the visual feedback.

For generating the reaching trajectories shown in Fig. 5 and 6, the starting position of the end effector is chosen such that it can be perceived by the stereo vision system. The direction vector $\Delta\theta$ is always calculated according to Eq. (5). The variable *step_size* is assigned with a value of 1mm to achieve the best reaching accuracy. Assigning a larger value to *step_size* magnifies the error in the extracted Jacobian $J$ and increases the curvature in the reaching trajectory as shown in Fig. 7. A larger *step_size* can cause the end effector to oscillate around the target before settling down. However, it also brings the benefit of reducing the requirement for the frequency of visual feedback and the computational burden associated with it. During the physical experiments on Nico, the *step_size* is initially set high and is reduced after the end effector has been moved close to the target. Fig. 8 shows a series of pictures capturing an actual reaching trajectory during a physical experiment on Nico. A wooden ball is attached to the distal end of Nico's arm as the end effector. A wooden ball of the same size is used as the reaching target. The centers of the end effector in the eight different positions on the trajectory are marked with white circles and superimposed onto the original pictures. The trajectory looks quite natural and is only slightly more curved than the trajectories recorded by Morasso in his 1981 paper. This larger curvature can be explained by the fact that Morasso used adults for his experiments. An adult's hand has a much larger range of motion compared with his/her head. So ignoring the translational effect of the neck joints has a lesser impact on the quality of the forward model and hence results in straighter reaching trajectories.

## 4. Discussion

The results of both simulations and physical experiments shown in the previous section provide support from the perspective of robotics for the feasibility of substituting neck proprioception with sensory information from the gyroscope/vestibular system during motor learning. In fact, the same substitution is used by humans for an important function - the vestibular-ocular reflex (VOR) (Kandel et al., 2000). VOR actively uses the sensory output of the vestibular nerve instead of the neck proprioception to adjust eye orientations to stabilize images on the retinas. To control eye movement, both angular velocity and angular position signals are required. The vestibular nerve outputs velocity signals only. These velocity signals are integrated in the brain stem to obtain position signals. The result of this integration can be directly used to transform the position of reaching target from the eye-centered coordinate system into the body-centered coordinate system. Since no individual-specific parameters are needed for this transformation, it does not have to be learned and can be hard-coded in the genes. All simulations and physical experiments previously described are based on Nico that is designed to match an average one-year-old. The naturally curved trajectories generated by our biologically plausible approach strongly suggest that a similar approach can be used by infants for learning to reach.

The previous section implicitly assumes that the body stays stationary relative to the outside world. If the assumption is violated, e.g. the body is passively rotated together with the head, $\theta_{gyro}$ registers a change, but the position and the orientation of the head relative to the body have not changed. Constructing the rotation matrix R in Eq. (3) according to the value of $\theta_{gyro}$ under this situation will lead to a reaching failure. However, this problem can be easily solved by allowing the integration of $\dot{\theta}_{gyro}$ only during the time the neck is moving.

According to (Buneo et al., 2002), there are two schemes which can be used to reach a target. The first scheme transforms the position of the target from the eye-centered coordinate system into the body-centered coordinate system before the reaching movement is initiated. The second scheme uses the target position in the eye-centered coordinate system directly for reaching. We have assumed that the former system is correct for our learning algorithms, but the latter assumption can also be easily

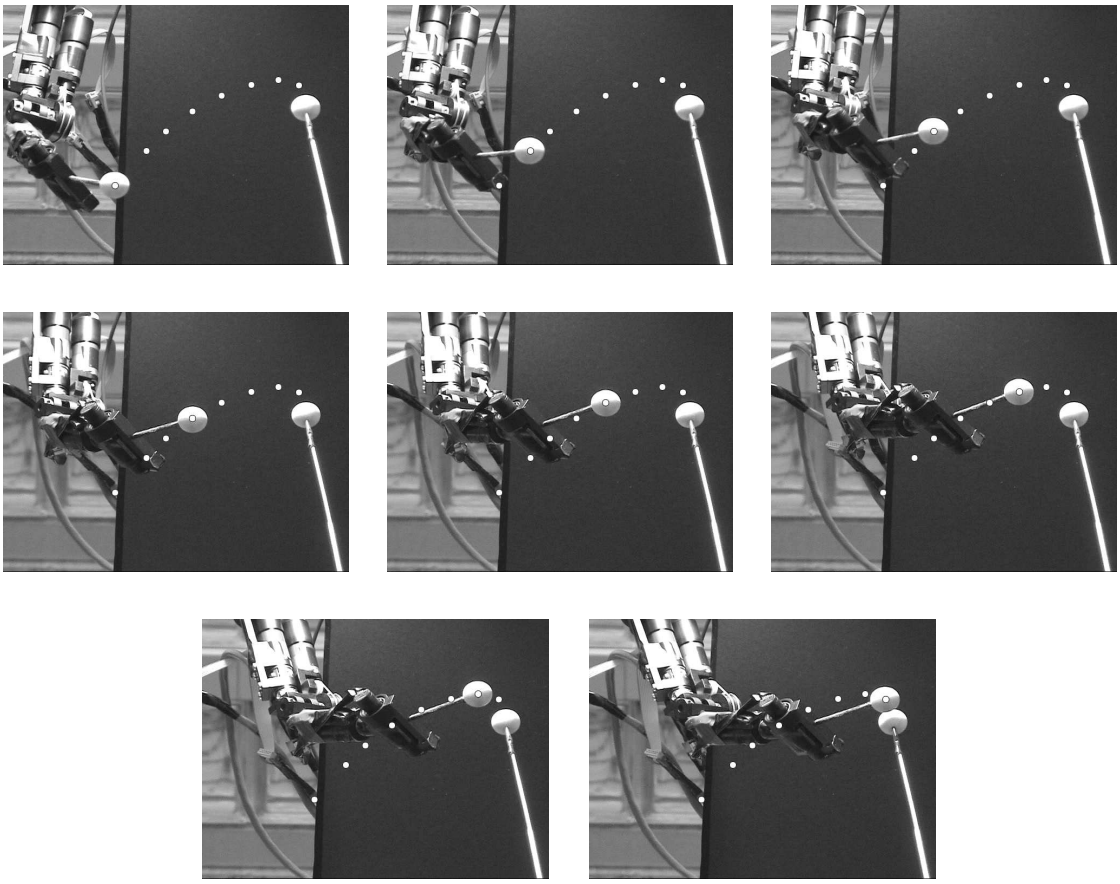

Figure 8: An actual reaching trajectory captured during a physical experiment on Nico.

implemented with the same architecture. The target position $x$ in the eye-centered coordinate system is first transformed into $x'$ as described in the previous section. Though in a strict sense, $x'$ is not exactly the target position in the body-center coordinate system, it can be used as a substitute for that without sacrificing much reaching accuracy. In order to be useful for the second scheme, our approach can be slightly modified, such that a forward model is learned with training samples in the form of $((\theta_{arm}, \theta_{gyro})_{input}, (x)_{output})$, where $\theta_{gyro}$ is a three dimensional vector representing the sensory readings of the gyroscope/vestibular system. After training, the target position in the eye-centered coordinate system can be directly used for reaching.

It is worth noting that if samples in the form of $((\theta_{arm}, \theta_{gyro})_{input}, (x)_{output})$ are used for learning the forward model, the dimension of the input is 9, only one less than that of the input of the original $f_{arm\_neck}$. While this seems to be only a small change in the dimensionality for our robot system, the advantage will scale with additional joints in the neck. Since the human neck has more than 4 DOFs even only at the joint level, this reduction can be substantially more. On the level of muscle activation, using $\theta_{gyro}$ to represent the principle effect of the head posture delivers enormous advantage for learning to reach.

One notable paper (D'Souza et al., 2001) studying sensorimotor learning on a robotic platform tries to solve the degrees-of-freedom problem with a sophisticated statistical learning algorithm called Locally Weighted Projection Regression (LWPR). LWPR is used to learn the direction mapping from $\delta x$ to $\delta\theta$ with the effort put into reducing the dimensionality locally. Our experiments suggest that the dimensionality of the problem can be reduced globally by substituting all DOFs in the neck altogether with sensory readings from the gyroscope, of which every human possess an equivalence, the vestibular system. Furthermore, our approach is based on a well-studied learning algorithm (RBFN) for which there are a large number of existing implementations.

## 5. Conclusion

In this paper we have demonstrated (1) that sensory information from a vestibular system can be used to replace a portion of the kinematic chain involved in learning to reach, (2) that this replacement can result in a decrease of the dimensionality of the learning problem, and (3) that the errors that are necessarily introduced by this replacement strategy lead

to curved reaching trajectories similar to those observed in human reaches. We have shown that these results hold true both in simulation and on a humanoid robot that is specifically designed to match an average one-year-old.

## References


Bernstein, N. A. (1967). *The co-ordination and regulation of movements*. Pergamon Press.

Bullock, D., Grossberg, S., and Guenther, F. H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, 5(4):408–435.

Buneo, C. A., Jarvis, M. R., Batista, A. P., and Andersen, R. A. (2002). Direct visuomotor transformations for reaching. *Nature*, 416:632–636.

D'Souza, A., Vijayakumar, S., and Schaal, S. (2001). Learning inverse kinematics. In *Proceedings of 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Flash, T. and Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, 5:1688–1703.

Guenther, F. H. and Barreca, D. M. (1997). Neural models for flexible control of redundant systems. In Morasso, P. G. and Sanguineti, V., (Eds.), *Self-organization, computational maps, and motor control*, pages 383–421. Elsevier.

Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394:780–784.

Hogan, N. (1984). An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, 4:2745–2754.

Jordan, M. I. and Rumelhart, D. (1992). Supervised learning with a distal teacher. *Cognitive Science*, 16:307–354.

Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (2000). *Principles of Neural Science*. McGraw-Hill, New York, fourth edition.

Mehta, B. and Schaal, S. (2002). Forward models in visuomotor control. *Journal of Neurophysiology*, 88(2):942–953.

Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, 42:223–227.

Pouget, A. and Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature Neuroscience*, 3:1192–1198.

Sun, G. and Scassellati, B. (2004). Reaching through learned forward model. In *Proceedings of 2004 IEEE/RSJ International Conference on Humanoid Robots*.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915.

Uno, Y., Kawato, M., and Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement: Minimum torque-change model. *Biological Cybernetics*, 61:89–101.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269:1880–1882.