Hybrid IDM/Impedance Learning in Human Movements

E Burdet, KP Tee, CM Chew, J Peters and V Loo BT
National University of Singapore, dept of Mech Eng
e.burdet@ieee.org, http://guppy.mpe.nus.edu.sg/~eburdet

Abstract

In spite of motor output variability and the delay in the sensori-motor, humans routinely perform intrinsically unstable tasks. The hybrid IDM/impedance learning controller presented in this paper enables skilful performance in strong stable and unstable environments. It considers motor output variability identified from experimental data, and contains two modules concurrently learning the endpoint force and impedance adapted to the environment. The simulations suggest how humans learn to skillfully perform intrinsically unstable tasks. Testable predictions are proposed.

1 Introduction

Problems arising when designing the control of a robot give precious information about the problems encountered by the CNS to control the limbs. However, the neuro-mechanical system of the human arm differs clearly from the robot hardware, in particular in the following two aspects:

1. *The CNS cannot use feedback to stabilize arm movements*; Neural feedback, whether operating through involuntary (reflex) or voluntary commands to our muscles, acts only with a delay of at least 60 ms [1]. In free movements, stability is insured by the muscle elastic property: when the hand is slightly disturbed it tends to return to the undisturbed trajectory [2, 3, 4, 5]. However, it is not clear how humans are able to perform stable movements in unstable environments characterized by negative impedance stronger than the arm impedance measured during free movements.

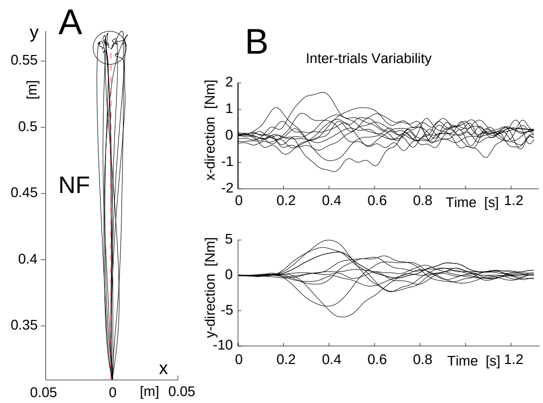

Figure 1: Variability in horizontal free arm movements measured in [7]. A: Ten movements performed in a null force field environment (NF). The coordinates are taken relative to the shoulder. B: Inter-trials variability identified from the movements of A.

2. *There is a large inter-trial variability in human movements*. The movements of robots are usually fairly reproducible with 0-mean noise when averaged over time [6], so that the trajectory remains around the planned trajectory. In contrast, when humans repeat point-to-point arm movements, most trials remain on either side of the mean trajectory. For example, in Fig.1A, most trials are either on the left or on the right of the straight line joining the start to the target and corresponding to the mean trajectory.

Engineers would doubt that a machine with such delayed feedback and motion variability is able to perform stable movements. Instability will amplify motor output variability and can lead to unsuccessful movements diverging in an unpredictable way from the planned trajectory. However, humans routinely perform intrinsically unstable tasks when using tools. When we lift a weight overhead the gravity will tend to let this weight fall to one side or the other. Similarly, material irregularities may perturb the chisel of a novice sculptor to the right or to the left of the intended path, yet an experienced artist has learned to skillfully compensate for such instability.

Previous work on learning in novel dynamic environments has shown that subjects learn an Internal Dynamic Model (IDM) of the task [8, 9, 10, 11]. The dynamics investigated in these studies all produced stable interactions between the arm and the environment. However in the real world many interactions are inherently unstable. [7] has observed how humans learn to stabilize movements performed in strongly unstable environments produced by a robotic interface.

We have investigated how humans can achieve skilful action in unstable dynamics by realizing a model of the adaptation occurring when movements are performed in arbitrary dynamics. In this paper we will describe this model, analyze its performance for movements repeated in typical stable and unstable force fields and formulate testable predictions.

[8, 9, 10] showed how the CNS can form an IDM of dynamics inducing a stable interaction with the arm. By repeating movements in unstable dynamics it may be possible to form an IDM. However, it is unlikely that this IDM could help in performing stable movements. Using the IDM to compensate for the environment instability would require the current position and velocity during movement. However, the significant delay of the sensori-motor loop and the unpredictability of motor output variability prevent accurate estimation of the position and velocity.

Therefore, we assume that the CNS has two modules to learn an IDM and the endpoint impedance. The CNS could modify the impedance at the endpoint of the arm to perform successful movements in unstable environments. A central question is how the CNS can coordinate the learning of an IDM and the learning of endpoint impedance. An error relative to the planned trajectory could require modification of either the IDM or the endpoint impedance, and it is not clear to which subsystem this error should be attributed. To control the concurrent learning of IDM and impedance, the hybrid controller analyzes how previous modifications affect the current control performance.

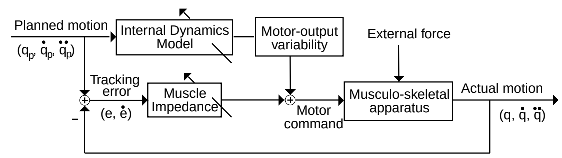

Figure 2: Control scheme of the neuro-mechanical control of human movements corresponding to Eq 5. Both the IDM and endpoint impedance are adapted to the task dynamics and the environment.

Algorithms to learn an IDM have been extensively investigated by the neural networks and control theory communities [12, 6], and successful robotic implementations have demonstrated the efficiency of these algorithms [6, 13]. In contrast, we found only a very few papers addressing impedance learning [14, 15, 16] and only one hybrid IDM/Impedance controller [17, 18] ([19, 20, 21] use the impedance control formulation to learn an IDM resulting in compliant motions). [14, 15, 16] use neural networks to learn optimal impedance and provide simulations showing smooth control in contact tasks and for [16] smooth transition from free movement to contact. [17, 18] modeled the control of the human arm using a Jordan type recursive neural network to input a Hill muscle model. However the very long learning phase and the too high reflexive components do not correspond to real experiments. To our knowledge, we present here the first controller able to stabilize unstable dynamics and considering a large motor output variability, similar to common human tasks.

2 Motor Control Model

During movements, the muscles have to produce force (denominated here by MUSCLE) corresponding to the arm dynamics (ARM) and to counteract environmental forces (FORCE). In equation form this gives:

$$\text{MUSCLE} = \text{ARM} - \text{FORCE}\,, \qquad (1)$$

where all these terms are expressed in the Cartesian space. The *mechanical impedance*, defined as the resistance to infinitesimal perturbations of the hand, is characterizing motion stability. Assuming that MUSCLE is a function of the position, velocity and activation $\mathbf{u}$, the force due to muscle impedance is equal to:

$$\text{IMP} = \mathbf{J}(\mathbf{q})^{-T} \left(\mathbf{K} \left(\mathbf{q}_p - \mathbf{q} \right) + \mathbf{D} \left(\dot{\mathbf{q}}_p - \dot{\mathbf{q}} \right) \right) , \quad (2)$$

where $\mathbf{J}(\mathbf{q})$ is the Jacobian of transformation between joint and Cartesian coordinates. $\mathbf{q}_p$ and $\mathbf{q}$ correspond to the planned and realized trajectories expressed in shoulder and elbow angles coordinates. $\mathbf{K}$ and $\mathbf{D}$ are the arm stiffness and damping expressed in the same coordinates:

$$\mathbf{K} = \frac{d\tau}{d\mathbf{q}} + \frac{d\tau}{d\mathbf{u}}\frac{d\mathbf{u}}{d\mathbf{q}} , \quad \mathbf{D} = \frac{d\tau}{d\dot{\mathbf{q}}} + \frac{d\tau}{d\mathbf{u}}\frac{d\mathbf{u}}{d\dot{\mathbf{q}}} . \quad (3)$$

Note that $\mathbf{K}$ and $\mathbf{D}$ include both intrinsic and reflexive components [22]. Equ. 2 corresponds to the linear term of MUSCLE when this function is linearized around the planned trajectory $\mathbf{q}_p$. MUSCLE - IMP is thus a function of this planned trajectory. We suppose that this function corresponds to the IDM of the task plus motor output variability modifying this planned dynamics, therefore:

$$\text{MUSCLE} = \text{IDM} + \text{NOISE} + \text{IMP}. \quad (4)$$

Combining Eqs 1 and 4 gives the control equation of the arm interacting with the environment:

$$\text{IDM} + \text{NOISE} + \text{IMP} = \text{ARM} - \text{FORCE}. \quad (5)$$

Fig.2 shows the control scheme corresponding to this equation. We observe that this scheme is similar to a nonlinear robot controller, with the IDM corresponding to a feedforward term and the IMP to linear feedback. However, IMP corresponds rather to intrinsic muscle properties than to neural feedback.

2.1 Identification of Variability

The "NOISE" in Eq 5 corresponds to motor output variability. We will identify this variability using experimental data of movements performed in a *null force field* (NF) environment [7]. We assume that motor output variability is 0 in mean over the trials, and that movements have the same motor output variability in all conditions. Therefore we can extract the variability from 50 force free or *null force field* (NF) trials (10 of which are plotted in Fig.1A) and use this variability to model movements performed in other force fields.

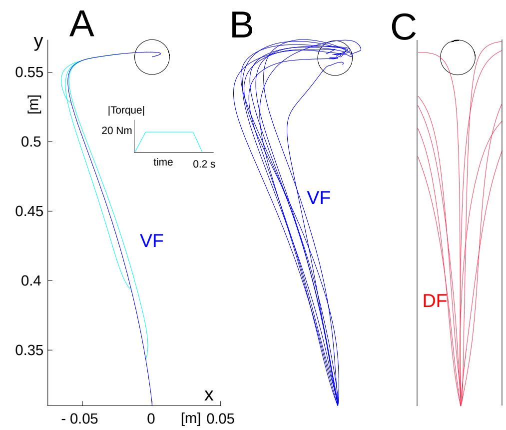

Figure 3: Stable and unstable movements. **A**: Perturbations exerted during movements in the VF without motor output variability. After every perturbation the trajectory tends to the undisturbed trajectory, thus the system of the arm interacting with the VF is stable. **B**: Before-effects in the VF, i.e. movements performed in the VF without learning. Most trajectories remain close to the planned trajectory, providing evidence that the arm interacting with the VF is stable. In contrast to the criterion of A, this criterion can also be used to show the stability of movements in real experiments with human subjects. **C**: Before-effects in the DF showing divergent thus unstable motions. The force field is shut down when the trajectory diverges from more that 0.03 m from the mean trajectory.

The arm model we use in this paper is a double pendulum with limbs of length (0.33, 0.34) m, center of mass (0.17, 0.19) m, mass (1.93, 1.52) kg and moment of inertia (0.0141, 0.0188) kg m^2 for the upper and lower arm respectively. The equation of the arm dynamics can be found in [5].

Supposing that all 50 trials use the same planned trajectory, we first identify this planned trajectory using the mean of Eq 5 over these 50 trials. We then use the same equation to identify the variability in every of the 50 trials (Fig.1B). We see that the variability increases with movement speed [23]. Furthermore the planned trajectory as defined by Eq 2 is superimposed on the mean trajectory (Fig.1A).

2.2 Stable and Unstable Interactions

To investigate learning of stable dynamics, we simulate movements in a *velocity dependent force field* (VF) similar to previous studies [8]. The force (in N) exerted on the hand is computed as a linear function of the hand speed $(\dot{x}, \dot{y})$ (m/s):

$$\text{FORCE} = \begin{bmatrix} 18 & -18 \\ 18 & 18 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} . \qquad (6)$$

The system of the arm interacting with the environment is *stable* when it is possible to make consecutive trials remain within an arbitrary small neighborhood by restricting the initial conditions to a correspondingly small neighborhood. Otherwise it is *unstable*. This corresponds to the definition of Lyapunov stability [24]. To test the stability of movements in the VF, we applied perturbations of small amplitude and duration during trajectories without noise (Fig.3A). If the system of the arm interacting with the environment is stable, the trajectory will return to the undisturbed trajectory after every perturbation. If the system is unstable the trajectory will diverge for some perturbation. According to this criterion the interaction with the VF is stable, as we can see in Fig.3A.

The above definition cannot be applied directly to experimental data due to motor output variability. However, when consecutive movements have very similar trajectories in spite of such variability, this strongly advocates for the stability of the arm interacting with the environment. Therefore, the trajectories with motor output variability performed in the VF of Fig.3B constitute strong evidence of stability.

On the other hand it is obvious that diverging trajectories correspond to an unstable system. In the *divergent force field* (DF) a force proportional to the distance from the planned trajectory is exerted on the hand during movement:

$$\text{FORCE} = \begin{bmatrix} 450\, x \\ 0 \end{bmatrix} . \qquad (7)$$

We can observe the unstable trajectories resulting from the interaction with the DF in Fig.3C. To avoid trajectories converging to infinity (virtually braking the arm!) the force field was shut down when the movement departed a 0.03 m safety zone around the planned trajectory.

3 Learning Algorithm

3.1 IDM Learning

Reforming Eq 5 and assuming that NOISE = 0, we have

$$\text{IMP} = (\text{ARM} - \text{FORCE}) - \text{IDM} \,, \qquad (8)$$

meaning that the impedance corresponds to the tasks dynamics (ARM - FORCE) not contained in the IDM. Therefore, IMP can be used as a teacher to learn novel dynamics in a supervised way. This approach has been used by [25, 26] and many others to learn the arm dynamics during movements, and further correspond to algorithms from nonlinear adaptive control successfully implemented on robots [6].

As our experiment will involve only a single movement, the IDM recorded along this movement can be memorized (as a function of time) in a look-up-table, corresponding to iterative learning control [27, 6]. We are currently developing a version of the hybrid IDM/Impedance controller based on a similar learning principle and valid in the whole workspace. To update the IDM, we use an iterative learning law modified from [6]:

$$\text{IDM}(i+1) = \text{IDM}(i) + \alpha\, \overline{\text{IMP}}(i) + \beta(i)\, \text{IMP}(i) \,, \qquad (9)$$

where i is the trial number and $\overline{\text{IMP}}(i)$ is the mean impedance over the previous trials. The term $\beta(i)\,\text{IMP}$, with a large but fast decreasing learning factor $\beta(i)$, ensures fast learning of novel dynamics. A potential problem is that this term brings motor output variability to the IDM. In the case of stable interactions the corresponding error will be bounded and have small amplitude. In the case of unstable interactions, however, this variability may be amplified, as in the DF, and lead to divergent trajectories. To prevent this we require that the $\beta(i)$ decreases fast. The $\alpha\,\overline{\text{IMP}}$ term, with a small constant α, filters motor output variability and insures that in the long term the mean dynamics will be compensated for by the IDM.

3.2 Impedance Learning

When the interaction with the environment is stable, the movement dynamics are reproducible and can be learned and well compensated for by the IDM, so the impedance term will soon become small due to IDM learning. When the motor output variability is high, the IDM learns to compensate for the mean dynamics, but cannot deal with unexpected variations of the motor command. Therefore the impedance has to be modified at the endpoint of the arm to make the control robust to motor output variability.

In this paper, we assume for simplicity that stiffness and damping matrices $\mathbf{K}$ and $\mathbf{D}$ are constant along the movement, although we have found similar results with time dependent impedance. We further assume that $\mathbf{D} = \mathbf{K}/5$.

The impedance learning algorithm is based on the observed behavior [7] that impedance increases to compensate for destabilization from the environment. Let the stiffness $\mathbf{K}$ be composed of two parts:

$$\mathbf{K} = \mathbf{K}_{IDM} + \mathbf{K}_S(i) \,. \tag{10}$$

$\mathbf{K}_{IDM}$ depends on the IDM, i.e. on the force to move the arm along the planned trajectory and on the force produced to compensate for the mean of external forces over the trials. For simplicity, we assume here that $\mathbf{K}_{IDM}$ is constant during the motion and equal to [40 20; 20 30] N m. $\mathbf{K}_S(i)$ is learned to prevent destabilization from the environment. After every trial, the stiffness $\Delta\mathbf{K}(i)$ to counteract destabilizing environmental forces can be identified using

$$\mathrm{FORCE}(i) - \overline{\mathrm{FORCE}}(i) = \Delta\mathbf{K}(i)\,\mathbf{e}(i) \,, \tag{11}$$

where $\mathbf{e}(i) = \mathbf{q}_p - \mathbf{q}$ is the tracking error in i-th trial and $\overline{\mathrm{FORCE}}(i)$ is the mean endpoint force over the trials. $\Delta\mathbf{K}(i)$ is corrupted by motor output variability. We prevent this variability from entering $\mathbf{K}_S$ by using the mean $\overline{\Delta\mathbf{K}}(i)$ over the trials and by learning $\mathbf{K}_S$ progressively: $\mathbf{K}_S(i)$ is updated according to:

$$\mathbf{K}_S(i+1) = \,(1-\lambda)\,\mathbf{K}_S(i) \,+\, \lambda\,\overline{\Delta\mathbf{K}}(i) \,, \tag{12}$$

where $1 > \lambda > 0$ is the learning factor. The force signal can be measured directly or be identified from the addition of the IDM and the impedance terms (fig.2). In the biological perspective, the tracking error $\mathbf{e}$ may be obtained from kinesthetic information provided by muscles spindles and the FORCE by Golgi tendons organs.

3.3 IDM-Impedance Coordination

The proposed learning algorithm is as followed. We assume that the arm has already learned the IDM for

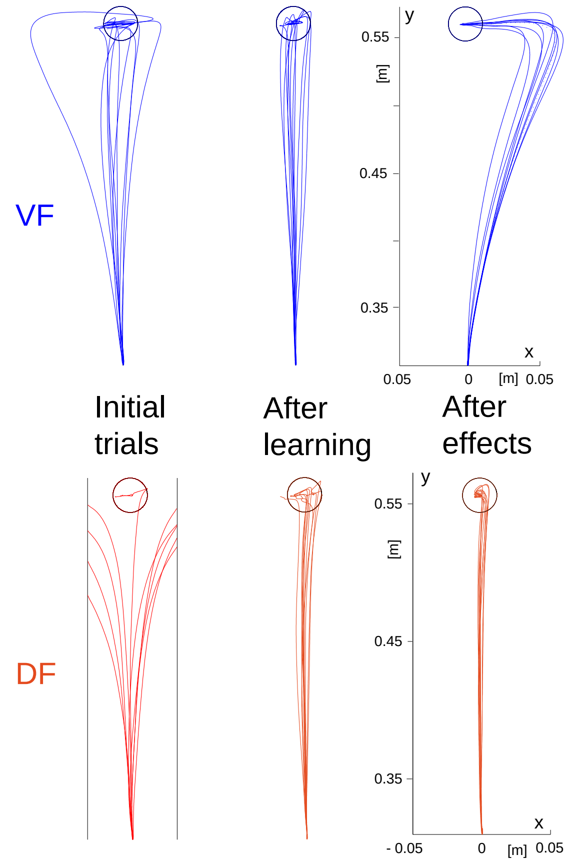

Figure 4: Evolution of trajectories during learning a stable interaction (VF) and an unstable interaction (DF).

movement in a null field environment. When the dynamics of the environment is changed and the actual movement deviates from the planned movement, this constitutes a *surprise*. The IDM learning is then activated to compensate for the new dynamics. In the consecutive trials, if the IDM learning is the correct strategy, the movement error will decrease. When the IDM learning doesn't reduce the error after several trials, this indicates that this strategy is not sufficient and the impedance learning will be activated.

In the simulation, the IDM is activated based on two indicators. They are the absolute error $\overline{ae}$ and the gradient of the least square fit of the absolute error $\overline{g}$, both computed over the last five trials. IDM learning is activated whenever $\overline{ae} > \delta_{\overline{ae}}$ and $\overline{g} > \delta_{\overline{g}}$, where $\delta_{\overline{ae}}$ AND $\delta_{\overline{g}}$ are some positive constants. $\delta_{\overline{ae}}$

and $\delta_{\overline{g}}$ are chosen such that this condition indicates that the tracking error is deteriorating. This condition is usually met when the movement is performed in a new dynamic environment which differs significantly from prior environment. If this condition is met six times in the last ten trials, the hybrid controller concludes that the IDM learning has not been successful, and it activates the impedance learning. Next section will show that this simple coordination strategy succeeds in learning stable and unstable dynamics.

4 Simulation Results

The experiment consists of learning the horizontal point-to-point movement of Fig.1A in the VF and DF defined by Eqs 6 and 7 respectively. The planned trajectory and the motor output variability were obtained from real data, by assuming that the variability is zero in mean over the trials (subsection 2.1). The IDM was initialized with the inverse dynamics to drive the arm along the planned trajectory. IDM learning used the learning factors $\alpha = 0.03$ and $\beta(i) = 0.8/i$, and the impedance learning factor was $\lambda = 0.03$. The thresholds to coordinate these two kinds of learning were $\delta_{\overline{ae}} = 0.001 m^2$ and $\delta_{\overline{g}} = 0.0007 m^2/\text{trial}$.

We analyze the learning using the *absolute error*

$$\int_0^T |x(t)| \, |\dot{y}(t)| \, dt \tag{13}$$

corresponding to the area between the actual movement and the planned path, and the *signed error*

$$\int_0^T x(t) \, |\dot{y}(t)| \, dt \, , \tag{14}$$

indicating the mean direction in which the path deviates from the planned path. The termination time T is determined using a curvature threshold of 0.07 m^{-1}.

We see in Fig.5 that in the VF the error decreases quickly and monotonously during learning. The *after effects*, movements performed after learning when the force field is cancelled, are about symmetric to the *before effects* (BE) or movements performed in the force field before, i.e. without learning. The absolute error also decreases in the DF, but the signed error remains in mean close to zero. The after effects have very small error and the signed error is not symmetric to the error for the before effects. This indicates that the hybrid controller uses distinct strategies to learn in stable and unstable dynamics.

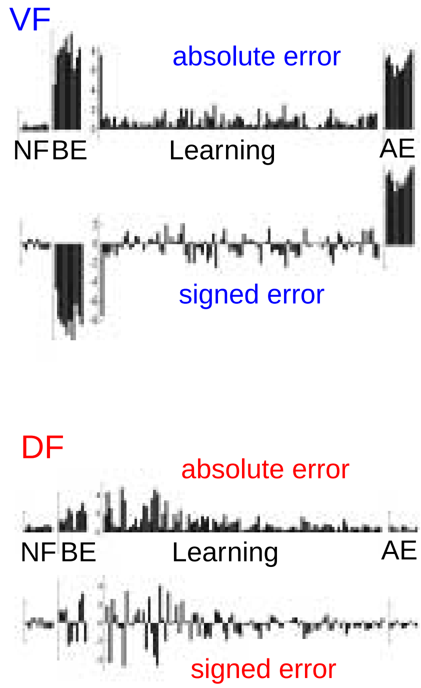

Figure 5: Evolution of path error during learning a stable interaction (the VF) and an unstable interaction (DF). NF stands for null fields movements, BE for before effects and AE for after effects, movements performed after learning when the force field is unexpectedly removed.

Fig.4 shows the ten initial trials, and ten trials performed after learning and in after-effects. In the VF, the trajectories, initially deformed to the left by the force field, converge rapidly to the straight path. Movements after learning are similar to NF movements, and after-effects are right from the straight line, corresponding to the inverse of the VF dynamics. This suggests that compensation for the VF is performed by the IDM. Fig.6 confirms that the IDM corresponds closely to the tasks dynamics.

Initial movements performed in the DF diverge widely to the left or to the right and are thus unstable. However, with learning the movements become

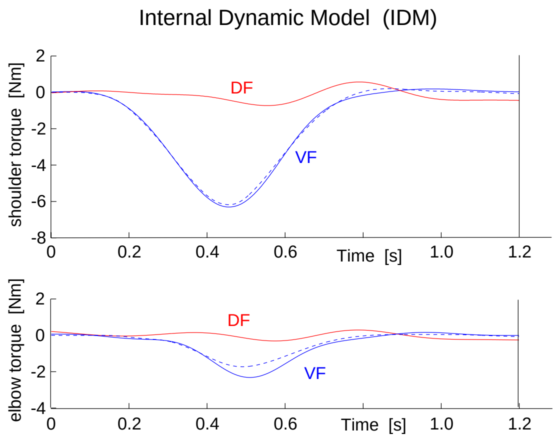

Figure 6: Internal Dynamic Model (IDM) learned during learning. Relative to the feedforward torque in the VF, the one in DF is of a small magnitude. The IDM learned in the VF (solid line) corresponds closely to the task dynamics in the VF (dashed line).

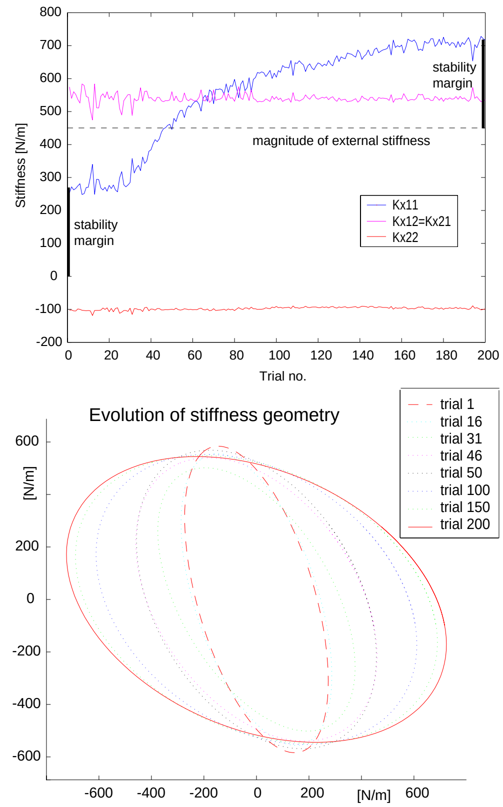

Figure 7: Evolution of stiffness during learning the DF. **A**: The Kx(1,1) parameter starts to increase around the 25th trial and converges before 200 trials. Kx(1,2) and Kx(2,2) remain constant. **B**: The effect of this evolution on the stiffness is seen by an elongation of the ellipse in the direction of the destabilizing force from the environment.

stable and similar to NF movements. After-effects are characterized by trajectories significantly closer to the straight line than NF movements, suggesting that the DF is not compensated for by the IDM. We observe in Fig.6 that in the DF the increase of IDM is negligible after learning.

How does the hybrid controller compensate for the DF? We can observe the modification of stiffness during learning the DF in Fig.7. While endpoint stiffness does not increase parallel to the movement and in the non-diagonal terms, stiffness increases gradually in a direction normal to the movement, to counteract destabilization from the DF. The stiffness ellipse, showing the force corresponding to a unit displacement, is gradually elongated in the direction of instability during learning. We note that arm stiffness is increased such that in the DF the total stiffness of the environment and the arm becomes similar to NF movements.

5 Discussion

The fact that deafferented primates can perform point-to-point well suggests that neural feedback is not necessary to control motion [28, 29]. This result was interpreted in mainly two different ways:

1. Kawato concluded that the CNS uses an IDM to compensate for the task dynamics. Sufficient evidence is now supporting this hypothesis [8, 9, 10, 11], but this does not explain how unstable tasks can be performed successfully.

2. Bizzi, Hogan and others concluded that movements are just planned kinematically and muscle impedance ensures that the arm moves along the planned trajectory, similar to the PD control strategy used by most industrial robots. Hogan further proposed that humans can voluntary control the impedance at the endpoint of their limbs to perform skillful movements [30]. However, measurement of impedance dur-

ing movement [5] suggested that if the CNS had to rely entirely on impedance to generate motions it would have to plan geometrically very complex trajectories in order to realize (geometrically very simple) point-to-point movements. Therefore [5] concluded that this control strategy is not plausible.

As a kind of synthesis of these two hypotheses, we propose that the CNS uses both IDM and Impedance learning to perform tasks in interaction with the environment. The hybrid IDM/Impedance learning controller introduced in this paper demonstrates the feasibility of this strategy. The simulations showed that this controller is able to perform well in stable and unstable dynamics. The hybrid IDM/Impedance controller combines fast IDM learning, leading to compensation of stable dynamics in only a very few trials, with slow impedance learning counteracting the effect of destabilizing dynamics.

While strong stability statements obviously require a theoretical analysis, the simulations performed in the VF and DF suggest that this controller is robust to a very high level of motor output variability and strong instability similar to the dynamics encountered in common tasks. We will extend this model by examining how muscles can produce endpoint force and impedance adapted to the dynamic environment, and by simulating simultaneous learning of several movements performed in the whole workspace.

Model's predictions can be tested in experiments measuring arm movements of human subjects in force fields produced by a robotic interface [7]. These predictions are:

- Motion stability in unstable environments is ensured by impedance learning.
- Impedance learning only takes place in environments producing an unstable interaction with the arm, while the Internal Dynamic Model learning is sufficient to compensate for stable dynamics.
- In stable interactions with the environment, the learning will be characterized by a fast, monotonous decrease of signed error relative to the planned or mean trajectory, and modifications of the IDM correlated with the novel dynamics.
- Learning unstable interactions is characterized by a non-monotonous decrease of signed error, and modifications of stiffness corresponding to the destabilization from the environment: If the environment is unstable in one direction only impedance will be increased specifically in this direction.

What is the ultimate goal of motor adaptation? In the algorithms of both the IDM and Impedance learning, muscle activation was modified to compensate exactly for external dynamics. This suggests that the goal of motor learning would be to enable the CNS to neglect the external influence, i.e. to make that the arm can move as if no force field was present. With motor adaptation, the CNS would not need to modify the hand-eye coordination or other complex sensorimotor coordination processes requiring a high level of computation and a very long learning. In analogy to physics laws arising from the invariance of some transformation, like the laws of special relativity constructed to be invariant under Lorentz transformations, the CNS would use a control strategy that is invariant under transformations of the dynamic environment.

Acknowledgments

We thank D.W. Franklin, M. Kawato, S. Keerthi, T.E. Milner, C.J. Ong and R. Osu for fruitful discussions (names in alphabetic order).

References

[1] P. Rack. Limitations of somatosensory feedback in control of posture and movement. In V.B. Brooks, editor, *Motor control. Handbook of Physiology*, volume Sect.1, pages 229–256. 1981.

[2] T.E. Milner and C. Cloutier. Compensation for mechanically unstable loading in voluntary wrist movement. *Experimental Brain Research*, 93:522–532, 1993.

[3] D.J. Bennett. Torques generated at the human elbow joint in response to constant position errors imposed during voluntary movements. *Experimental Brain Research*, 95:488–498, 1993.

[4] J. Won and N. Hogan. Stability properties of human reaching movements. *Experimental Brain Research*, 107:125–136, 1995.

[5] H. Gomi and M. Kawato. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biological Cybernetics*, 76(3):163–171, 1997.

[6] E. Burdet, A. Codourey, and L. Rey. Experimental evaluation of nonlinear adaptive controllers. *IEEE Control Systems Magazine*, pages 39–47, April 1998.

[7] E. Burdet, R. Osu, D.W. Franklin, T.E. Milner, and M. Kawato. The cns skillfully stabilizes unstable dynamics by learning optimal impedance. *(submitted)*.

[8] R. Shadmehr and F.A. Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5):3208–3224, 1994.

[9] J. Lackner and P. DiZio. Rapid adaptations to coriolis force perturbations of arm trajectory. *Journal of Neurophysiology*, 72:299–313, 1994.

[10] J.W. Krakauer, M.F. Ghilardi, and C. Ghez. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neuroscience*, 2:1026–1031, 1999.

[11] M. Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6):718–727, 1999.

[12] F.L. Lewis, C.T. Abdallah, and D.M. Dawson. *Control of Robot Manipulators.* Macmillan, 1993.

[13] S. Vijayakumar and S. Schaal. Fast and efficient incremental learning for high dimensional movement systems. In *IEEE International Conference on Robotics and Automation*, 2000.

[14] H. Asada. Teaching and learning of compliance using neural nets: representation of generation of nonlinear compliance. In *IEEE International Conference on Robotics and Automation*, pages 1237–1244, 1990.

[15] M. Cohen and T. Flash. Learning impedance parameters for robot control using an associative search network. *IEEE Transactions on Robotics and Automation*, 1991.

[16] T. Tsuji, K. Ito, and P.G. Morasso. Neural network learning of robot arm impedance in operational space. *IEEE Transactions on Systems, Man, Cybernetics -Part B: Cybernetics*, 26(2):290–298, 1996.

[17] S. Stroeve. Impedance characteristics of a neuromusculoskeletal model of the human arm i. posture control. *Biological Cybernetics*, 81(5/6):475–494, 1999.

[18] S. Stroeve. Impedance characteristics of a neuromusculoskeletal model of the human arm ii. movement control. *Biological Cybernetics*, 81(5/6):495–504, 1999.

[19] Y. Maeda and H. Kano. Learning control for impedance controlled manipulator. In *IEEE Conference on Decision and Control*, pages 3135–3140, 1992.

[20] H. Gomi and M. Kawato. Neural network control for a closed loop system using feedback-error-learning. *Neural Networks*, 6(7):933–946, 1993.

[21] C.C. Cheah and D. Wang. Learning impedance control for robotic manipulators. *IEEE Transactions on Robotics and Automation*, 14(3):452–465, 1998.

[22] E. Burdet, R. Osu, D.W. Franklin, T. Yoshioka, T.E. Milner, and M. Kawato. A method for measuring endpoint stiffness during multi-joint arm movements. *Journal of Biomechanics*, 33(12):1705–1709, 2000.

[23] R.A. Schmidt. Motor output variability: A theory for the accuracy of rapid motor acts. *Psychological Review*, 86(5):415–451, 1979.

[24] J.-J.E. Slotine and W. Li. *Applied Nonlinear Control*. Prentice-Hall, 1991.

[25] J.S. Albus. A new approach to manipulator control: The cerebellar model articulation controller. *Journal of Dynamical Systems, Measurements, and Control*, 97:220–227, 1975.

[26] M. Kawato, K. Furukawa, and R. Suzuki. A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics*, 57:169–185, 1988.

[27] S. Arimoto. Learning control. In M.W. Spong, F.L. Lewis, and C.T. Abdallah, editors, *Robot Control*, pages 1064–1069. 1992.

[28] E. Bizzi, N. Accornero, W. Chapple, and N. Hogan. Arm trajectory formation. *Experimental Brain Research*, 46:2738–2744, 1982.

[29] C. Ghez, J. Gordon, M.F. Ghilardi, C.N. Christakos, and S.E. Cooper. Role of proprioceptive input in the programming of arm trajectories. In *Cold Spring Harbor Symposia on Quantitative Biology LV*, pages 837–847, 1990.

[30] N. Hogan. The mechanics of multi-joint posture and movement control, biological cybernetics. *Biological Cybernetics*, 52:315–333, 1985.