Supplementary materials for: "Adaptation to inversion of the visual field: a new twist on an old problem."

Timothy P. Lillicrap[1], Pablo Moreno-Briseño[2], Rosalinda Diaz[2], Douglas B. Tweed[3], Nikolaus F. Troje[5,6], Juan Fernández-Ruiz[2,4]

1. Centre for Neuroscience Studies, Queen's University
2. Laboratorio de Neuropsicologìa, Facultad de Medicina, Universidad Nacional Autònoma de Mèxico.
3. Department of Physiology, University of Toronto
4. Facultad de Psicologìa, Universidad Veracruzana
5. Department of Psychology, Queen's University
6. School of Computing, Queen's University

March 3, 2013

Supplementary Equations:

The control problem studied here is schematized in Supplementary Figure 1 and works as follows: the agent is presented with a visual target presented on a fixed 2D plane. Given the perceived location of the target the agent must then try to touch it. To keep the simulations simple, the attempt to strike the target is open loop (i.e. the agent is not able to modify their choice in mid-flight) and happens in a single time step or trial. Let, $\mathbf{x}(t) \in \mathbb{R}^2$ and $\mathbf{x}^{\mathrm{p}}(t) \in \mathbb{R}^2$, be the real and perceived location of the target on trial t. Here we will denote the elements of a vector using subscripts, e.g. $\mathbf{x}_1(t)$ is the first element of $\mathbf{x}$ at trial t, and for convenience we will often drop the trial index, e.g. we may just write $\mathbf{x}$ or $\mathbf{x}^{\mathrm{p}}$. The target is viewed through one of three transformations, which are induced by prisms in the experimental setup,

1. No visual perturbation; the identity transformation:
   $\mathbf{x}^{\mathrm{p}} = T_{\mathrm{I}}(\mathbf{x}) = \mathbf{x}$

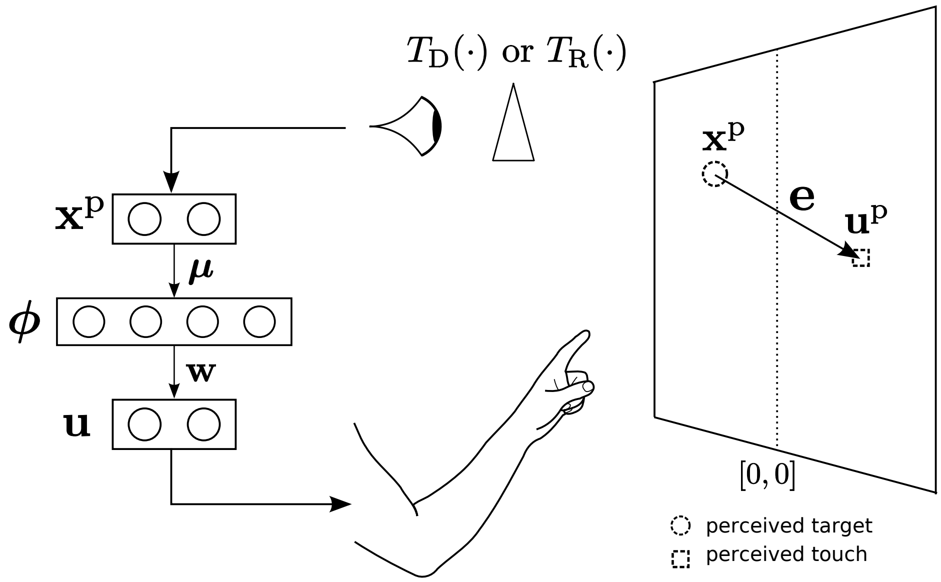

Figure 1: A schematic of the task and modelled controller. Each trial, the controller gets a perceived target, $\mathbf{x}^{\text{p}}$, as input and computes a command, $\mathbf{u}$, which dictates the position on a screen that the arm will touch. The feedback error, $\mathbf{e}$, associated with a trial is computed from the difference between the perceived target, $\mathbf{x}^{\text{p}}$, and perceived touch location, $\mathbf{u}^{\text{p}}$. Both the perceived target and touch location are viewed normally during baseline and washout conditions, and by way of reversing prism transformation, $T_{\text{R}}(\cdot)$, or displacing prism transformation, $T_{\text{D}}(\cdot)$, during perturbed conditions. The controller is composed of a vector of local basis functions, $\boldsymbol{\phi}$, and two sets of adjustable parameters, $\boldsymbol{\mu}$ and $\mathbf{w}$.

2. A horizontal displacing perturbation of magnitude D:
   $\mathbf{x}^{\text{p}} = T_{\text{D}}(\mathbf{x}) = [\mathbf{x}_1 + D, \mathbf{x}_2]$

3. A horizontal reversing perturbation about 0:
   $\mathbf{x}^{\text{p}} = T_{\text{R}}(\mathbf{x}) = [(-1) \cdot \mathbf{x}_1, \mathbf{x}_2]$

Let, $\mathbf{u}(t) = \mathbf{u}(\mathbf{x}^{\text{p}}(t)) \in \mathbb{R}^2$, be the location in real world coordinates that the agent touches and, $\mathbf{u}^{\text{p}}(t) = \mathbf{u}^{\text{p}}(\mathbf{x}^{\text{p}}(t)) \in \mathbb{R}^2$, be the perceived location of the agent's throw. The same transformation, $T_{\text{X}}(\cdot)$, is applied to the real world coordinates of the agent's throw to get the perceived location of the touch: $\mathbf{u}^{\text{p}}(\mathbf{x}^{\text{p}}(t)) = T_{\text{X}}(\mathbf{u}(\mathbf{x}^{\text{p}}(t)))$. If, as in the main text, we take the control law to be parameterized by the vector, $\boldsymbol{\alpha}$, then we can write the real world and perceived touch locations as, $\mathbf{u}(\mathbf{x}^{\mathrm{p}}(t);\boldsymbol{\alpha})$, and, $\mathbf{u}^{\mathrm{p}}(\mathbf{x}^{\mathrm{p}}(t);\boldsymbol{\alpha})$, respectively.

We define the error made on a given trial, t, to be the vector difference between the perceived target and touched location, $\mathbf{e}(t) = \mathbf{u}^{\mathrm{p}}(t) - \mathbf{x}^{\mathrm{p}}(t)$. Now, we define the loss function to be minimized as, $L = \frac{1}{2}||\mathbf{e}||^2$.

Having defined the loss function we are interested in minimizing, here we briefly demonstrate the effect of the various transformations on the partial derivatives of the error, $\mathbf{e}$, with respect to the commands, $\mathbf{u}$. We start by recalling the standard sensorimotor learning rule, which follows the gradient of the loss function with respect to the control parameters, given by:

$$\frac{\partial L}{\partial \boldsymbol{\alpha}_n} = \mathbf{e}^{\mathrm{T}} \frac{\partial \mathbf{e}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \boldsymbol{\alpha}_n} = \sum_{i=1}^{2} \mathbf{e}_i \cdot \sum_{j=1}^{2} \frac{\partial \mathbf{e}_i}{\partial \mathbf{u}_j} \cdot \frac{\partial \mathbf{u}_j}{\partial \boldsymbol{\alpha}_n} \tag{1}$$

where, $\boldsymbol{\alpha}_n$, is the n^{th} element of the parameter vector. Thus the standard update rule is:

$$\Delta \boldsymbol{\alpha}_n = -\eta \frac{\partial L}{\partial \boldsymbol{\alpha}_n} = -\eta \mathbf{e}^{\mathrm{T}} \frac{\partial \mathbf{e}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \boldsymbol{\alpha}_n} \tag{2}$$

where, $\eta > 0$, is a learning constant. Here we will examine the sensitivity matrix, $\partial \mathbf{e}/\partial \mathbf{u}$, under four visual perturbations. The first three are the ones used in experiments, and the fourth illustrates how a compressing/expanding prism would effect the sensitivity though no such prism was used.

1. Under the identity transformation, $T_{\mathrm{I}}(\cdot)$, we have: $\mathbf{e} = \mathbf{u}^{\mathrm{p}} - \mathbf{x}^{\mathrm{p}} = T_{\mathrm{I}}(\mathbf{u}) - \mathbf{x}^{\mathrm{p}} = \mathbf{u} - \mathbf{x}^{\mathrm{p}}$. Thus the matrix of partial derivatives is:

$$\frac{\partial \mathbf{e}}{\partial \mathbf{u}} = \frac{\partial \mathbf{u}}{\partial \mathbf{u}} - \frac{\partial \mathbf{x}^{\mathrm{p}}}{\partial \mathbf{u}} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \tag{3}$$

2. Under the displacing transformation, $T_{\mathrm{D}}(\cdot)$, we have: $\mathbf{e} = \mathbf{u}^{\mathrm{p}} - \mathbf{x}^{\mathrm{p}} = T_{\mathrm{D}}(\mathbf{u}) - \mathbf{x}^{\mathrm{p}} = \mathbf{u} + [\mathrm{D}, 0] - \mathbf{x}^{\mathrm{p}}$. Thus the matrix of partial derivatives is:

$$\frac{\partial \mathbf{e}}{\partial \mathbf{u}} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \tag{4}$$

   Notice that the displacing perturbation has no effect at all on the sensitivity matrix, which means that the standard update rule will function just as well as it did with no perturbation.

3. Under the reversing transformation, $T_{\mathrm{R}}(\cdot)$, we have: $\mathbf{e} = \mathbf{u}^{\mathrm{p}} - \mathbf{x}^{\mathrm{p}} = T_{\mathrm{R}}(\mathbf{u}) - \mathbf{x}^{\mathrm{p}} = [(-1) \cdot \mathbf{u}_1, \mathbf{u}_2] - \mathbf{x}^{\mathrm{p}}$. Thus the matrix of partial derivatives is:

$$\frac{\partial \mathbf{e}}{\partial \mathbf{u}} = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \tag{5}$$

It is important to notice that the sensitivity matrix under the reversing transformation, $T_{\mathrm{R}}(\cdot)$, has an element whose sign is opposite from the normal conditions (i.e. under the identity transformation). In particular, the sign of $\partial \mathbf{e}_1 / \partial \mathbf{u}_1$ has changed from $1 \rightarrow -1$. This means that the update rule given in equation 2 will be wrong under the reversing condition if the old matrix $\partial \mathbf{e}/\partial \mathbf{u}$ associated with the identity condition is used during learning.

4. Now, we consider a horizontal compressing/expanding transformation of magnitude, $\mathrm{C} > 0$, which acts like: $\mathbf{x}^{\mathrm{p}} = T_{\mathrm{C}}(\mathbf{x}) = [\mathrm{C} \cdot \mathbf{x}_1, \mathbf{x}_2]$. Under this transformation, we have: $\mathbf{e} = \mathbf{u}^{\mathrm{p}} - \mathbf{x}^{\mathrm{p}} = T_{\mathrm{C}}(\mathbf{u}) - \mathbf{x}^{\mathrm{p}} = [\mathrm{C} \cdot \mathbf{u}_1, \mathbf{u}_2] - \mathbf{x}^{\mathrm{p}}$. Thus the matrix of partial derivatives is:

$$\frac{\partial \mathbf{e}}{\partial \mathbf{u}} = \begin{pmatrix} C & 0 \\ 0 & 1 \end{pmatrix} \tag{6}$$

Since, none of the signs of the sensitivity matrix are altered by compresion ($0 < \mathrm{C} < 1$) nor expansion ($1 < \mathrm{C}$), the standard learning rule will continue to function. However, if the compression factor is very small, or expansion factor very large, it is possible that the learning rate, $\eta > 0$, may need to be adjusted for learning to procede nicely (i.e. quickly, and without instability).

5. Finally, we consider a transformation which is a rotation by some angle, θ, which acts like: $\mathbf{x}^{\mathrm{p}} = T_{\mathrm{Rot}(\theta)}(\mathbf{x}) = \mathrm{Rot}(\theta) \cdot [\mathbf{x}_1, \mathbf{x}_2]$. Where,

$$\mathrm{Rot}(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \tag{7}$$

Under this transformation, we have: $\mathbf{e} = \mathbf{u}^{\mathrm{p}} - \mathbf{x}^{\mathrm{p}} = T_{\mathrm{Rot}(\theta)}(\mathbf{u}) - \mathbf{x}^{\mathrm{p}} = [\cos(\theta)\mathbf{u}_1 - \sin(\theta)\mathbf{u}_2, \sin(\theta)\mathbf{u}_1 + \cos(\theta)\mathbf{u}_2] - \mathbf{x}^{\mathrm{p}}$. Thus the matrix of partial derivatives is:

$$\frac{\partial \mathbf{e}}{\partial \mathbf{u}} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \tag{8}$$

Thus, the signs of the matrix elements are dependent on the rotation angle, θ.

A Caveat: In the paper we said that a change in the sign of any of the elements of the normal sensitivity matrix will doom learning. This is usually, but not strictly true. It's actually the vector of partial derivatives $\partial L/\partial \mathbf{u}$ which determines whether learning will work or not. One can construct special conditions under which multiple sign changes in the sensitivity matrix occur, but the vector of partials $\partial L/\partial \mathbf{u}$ maintain the correct signs for learning. These are special cases and do not seem to have much bearing on most sensorimotor problems though; see [1] for a more detailed account of this.

Simulation details: In our simulations, the parameterized controller was a simple 1 hidden layer radial basis function network [8]. None of the ideas central to this paper should change for any other reasonable choice of controller. The network is schematized in Supplementary Figure 1. This type of model, in which the basis functions, $\boldsymbol{\phi}$, are local, was selected because it is widely assumed that visuomotor learning generalizes locally. The controller has the following form:

$$\mathbf{u}(\mathbf{x}^{\mathrm{p}}) = \left[\; \sum_{i=1}^{n} \mathbf{w}_{1i} \cdot \boldsymbol{\phi}_i(\mathbf{x}^{\mathrm{p}}), \;\; \sum_{i=1}^{n} \mathbf{w}_{2i} \cdot \boldsymbol{\phi}_i(\mathbf{x}^{\mathrm{p}}) \; \right] \tag{9}$$

where, $\mathbf{w} \in \mathbb{R}^{2\times n}$, is a weight matrix and the basis functions are locally tuned gaussians, $\boldsymbol{\phi}_i(\cdot), i \in \{1, ..., n\}$, given by:

$$\boldsymbol{\phi}_i(\mathbf{x}) = \boldsymbol{\beta}_i \cdot e^{\frac{-||\mathbf{x}-\boldsymbol{\mu}_i||^2}{2\boldsymbol{\sigma}_i^2}} \tag{10}$$

where, $\boldsymbol{\beta}, \boldsymbol{\sigma} \in \mathbb{R}^2$, and, $\boldsymbol{\mu} \in \mathbb{R}^{2\times n}$, are model parameters. In our simulations the vector, $\boldsymbol{\sigma}$, was hand tuned and fixed.

For convenience we can write the parameters flattened into a generic parameter vector, $\boldsymbol{\alpha} = \{\boldsymbol{\beta}, \boldsymbol{\mu}, \mathbf{w}\}$. Then the update rule we used during simulations is very similar to the standard rule given earlier; the only difference is the inclusion of a momentum term to speed learning:

$$\Delta\boldsymbol{\alpha}_n(t) = -\eta \frac{\partial L}{\partial \boldsymbol{\alpha}_n} = -\eta \mathbf{e}^{\mathrm{T}} \frac{\partial \mathbf{e}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \boldsymbol{\alpha}_n} + \rho \left[\Delta\boldsymbol{\alpha}_n(t-1)\right] \tag{11}$$

where, ρ, is the momentum term which speeds up learning by incorporating a weighted version of the previous trial's update [6].

Simulation overview: Two variants of the model performed essentially the same task as human subjects in the random-target experiments. Each trial, a target was chosen at random and the transformed version of this target was given to the model variants as input. Each model variant computed a command output which was transformed via the same perturbation after a small amount of Gaussian noise ($\mu = 0, \sigma = 1$) was added to the motor command. We performed n=15 runs of the simulation (the same as the number of subjects who participated in the random-target experiment) with different seeds for the random number generator. In a pre-experiment phase, both of the model variants were trained to correctly map perceived targets to the motor commands required to hit them without any visual transformation. Since initially the state of the network was randomly initialized, it took approximately 5000 trials to converge to a robust controller for the identity transformation. Once this phase was complete (i.e. performance is acceptable), the model variants were programmed to perform the reversal task with random targets.
For details about how the sensitivity is learned in the case of the implicit supervision variant, please see [1].

Supplementary Discussion:

Comments on "explicit" strategies: To reiterate, a common approach in sensory motor neuroscience is to attribute these sorts of unexplained results (i.e. erratic learning with no aftereffect) to "cognitive" or "explicit" strategies or "tricks". But these words are too vague to provide a real theory. We would do better to focus on the algorithmic properties we need to cover the facts, namely that the mechanism invoked by reversals is erratic, apparently lapsing or overcorrecting from one trial to the next, and it leaves no trace when sensitivity derivatives are restored to normal.

One could easily devise a learning mechanism with these two properties, but that approach would be ad hoc it could be done in many different ways and would give us no idea why natural selection favored this type of learning. A stronger theory would have to start from some algorithm that had already been identified as useful, based on fundamental control issues, and that showed the properties in question. There appear to be at least two promising approaches. Koerding et al. have argued on Bayesian grounds that learning should proceed on multiple timescales, some mechanisms learning quickly and forgetting quickly, others learning and forgetting slowly [5]. The rationale is that some changes in the world are fleeting and others long-lasting, and the brain should be able to cope with both. Our subjects lack of aftereffect, and perhaps their erratic progress, may be explained by very short-term learning. Another approach focuses on computational resources: Fortney and Tweed [3] have shown that complex tasks can be learned quickly and with few neurons by a mechanism called weightless learning, but the price is that the process is sometimes erratic and has no memory.

Comments on aftereffects: Negative aftereffects are a robust phenomenon that have been found in many conditions not only in the prism adaptation literature, but in a variety of sensorimotor adaptation tasks across both kinematic and dynamic domains [11, 4]. Negative aftereffects are thought to be a manifestation of the modification of an internal model in response to changes either in the environment or to changes in the physical plant [10]. Aftereffects are evident once the normal (or previous) conditions are reestablished. The absence of aftereffects has principally been observed in two experimental conditions. The first is when the internal model is simply not updated; the second is when the system learns to switch between two previously learned internal models based on context. [7, 9, 10].

Here we found typical negative aftereffects [2] in the fixed target refracting prism experiment; raw results can be seen in Supplementary Figure 2. The negative aftereffect is readily apparent when compared to Figure 1c of the main text where the absolute values of the error were plotted. The most intriguing findings, however, were the results obtained in the reversing prism experiments. The reversing prism group as a whole not only did not show a negative aftereffect, but seem to have a large positive aftereffect that disappears almost immediately followed by more subtle, but still positive, aftereffects.

The reversal experiments where the target appeared in random positions also shed light on the aftereffects results. Once the reversing prism was withdrawn, the group showed a large aftereffect that did not returned to baseline levels during the fifty last trials after the prism was withdrawn. We think those aftereffects were in fact negative aftereffects derived from the gradual maladaptive learning done while performing with the reversing prisms (i.e. an after effect stemming from learning with the wrong sensitivity estimate).

References

[1] M.N. Abdelghani, T.P. Lillicrap, and D.B. Tweed. Sensitivity derivatives for flexible sensorimotor learning. *Neural Computation*, 20(8):2085–2111, 2008.

[2] J. Fernández-Ruiz and R. Díaz. Prism adaptation and aftereffect: specifying the properties of a procedural memory system. *Learning & Memory*, 6(1):47–53, 1999.

[3] K.P. Fortney and D.B. Tweed. Learning without synaptic change: a mechanism for sensorimotor control. *COSYNE conference on Computational and Systems Neuroscience*, Salt Lake City:Abstract 104., 2006.

[4] F. Gandolfo, F.A. Mussa-Ivaldi, and E. Bizzi. Motor learning by field approximation. *Proc. Natl. Acad. Sci. U.S.A.*, 93(9):3843–3846, 1996.

[5] K.P. Körding, J.B. Tenenbaum, and Shadmehr R. The dynamics of memory as a consequence of op¡timal adaptation to a changing body. *Nature Neuroscience*, 10:779–786, 2007.

[6] Yann LeCun, Leon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. Efficient backprop. In Genevieve B. Orr and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade*. Springer, 1998.

[7] T.A. Martin, J.G. Keating, H.P. Goodkin, A.J. Bastian, and W.T. Thach. Throwing while looking through prisms. i. focal olivocerebellar lesions impair adaptation. *Brain*, 119(4):1183–1198, 1996.

[8] John Moody and Christian J. Darken. Fast learning in networks of locally-tuned processing units. *Neural Computation*, 1(2):281–294, 1989.

[9] G.M. Redding, Y. Rossetti, and B. Wallace. Applications of prism adaptation: a tutorial in theory and method. *Neuroscience and Biobehavioral Reviews*, 29(3):431–444, 2005.

[10] R. Shadmehr and S.P. Wise. Motor learning and memory for reaching and pointing. In M.S. Gazzaniga, editor, *The Cognitive Neurosciences*, pages 511–524. MIT Press, third edition edition, 2004.

[11] M.J. Weiner, Hallett. M., and H.H. Funkenstein. Adaptation to lateral displacement of vision in patients with lesions of the central nervous system. *Neurology*, 33(6):766–772, 1983.

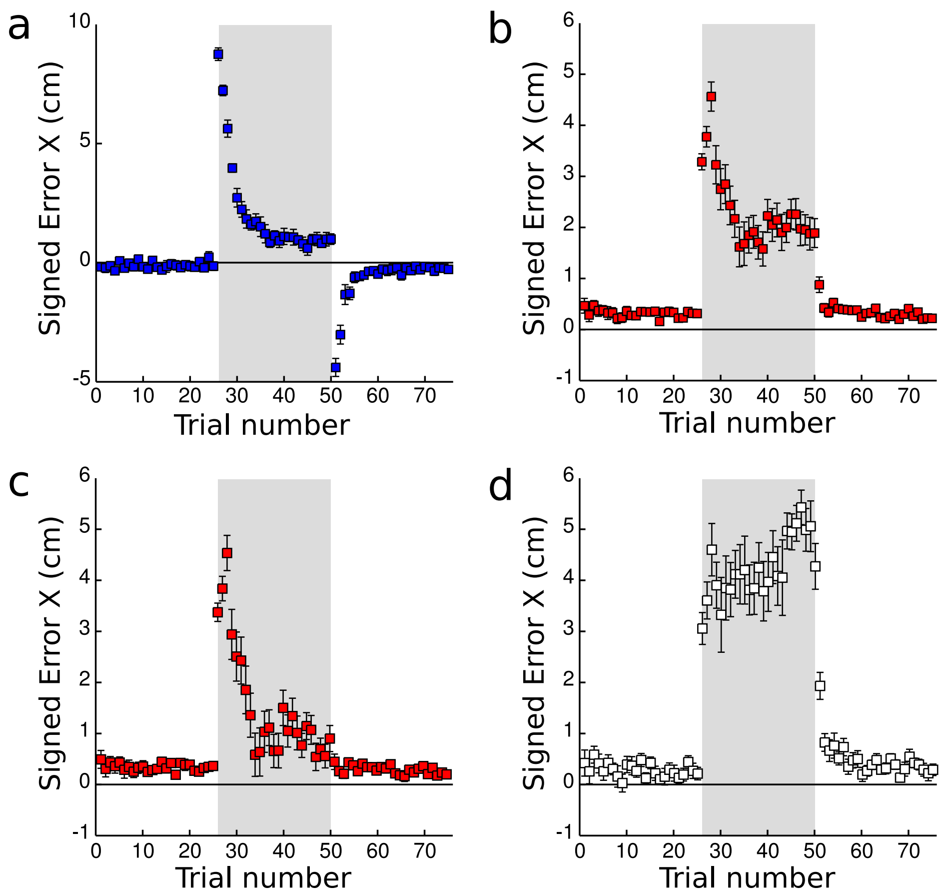

Figure 2: Affereffects in the fixed target experiments. a, Mean error (note, that this is *not* absolute error) made by subjects in the refracting prisms condition. Notice that the aftereffect in the washout phase is negative. b, Mean error made by subjects in the reversing prism condition. Notice that the aftereffect is small and positive. c, Mean error for the 55 of 78 subjects who were able to improve during the reversing prism condition (i.e. their mean absolute error during the last 5 perturbed trials was less than 3.0 cm). Notice that in this case there is no discernable afftereffect at all. d, Mean error for the 23 of 78 subjects who were unable improve during the reversing prism condition. Notice that in this case there is a easily seen positive aftereffect.

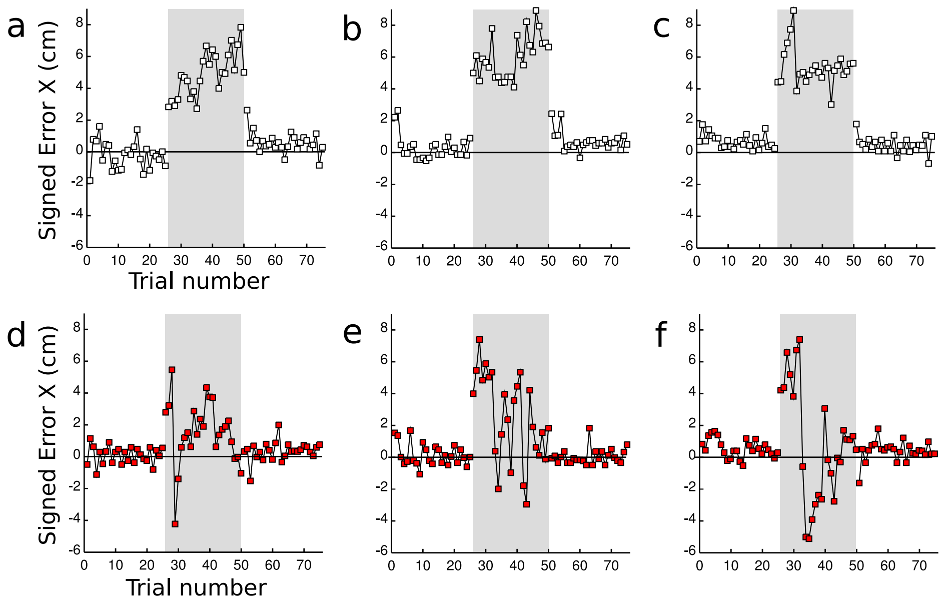

Figure 3: Examples of individual subject performance in the fixed target / reversing prism experiment. a-c, Three subjects who were unable to reduce the mean absolute error across the last 5 perturbed trials to less than 3 cm. Notice that these subjects show a tendency to correct in the wrong direction throughout the 25 perturbed trials. d-f, Shows the performance of three typical subjects who were able to reduce the mean absolute error across the last 5 perturbed trials to < 3 cm. However, these subjects are very unstable around the target; i.e. they do not exhibit the nice learning curves observed under the displacing prism condition.