Modular decomposition in visuomotor learning

Zoubin Ghahramani*† and Daniel M. Wolpert‡

†Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139, USA

‡Sobell Department of Neurophysiology
Institute of Neurology, Queen Square
London WC1N 3BG, United Kingdom

Nature **386**:392-395 (1997)

The principle of 'divide-and-conquer,' the decomposition of a complex task into simpler subtasks each learned by a separate module, has been proposed as a computational strategy during learning [1, 2, 3]. We explore the possibility that the human motor system uses such a modular decomposition strategy to learn the visuomotor map, the relationship between visual inputs and motor outputs. Using a virtual reality system, subjects were exposed to opposite prism-like visuomotor remappings—discrepancies between actual and visually perceived hand locations—for movements starting from two distinct locations. Despite this conflicting pairing between visual and motor space, subjects learned the two starting-point-dependent visuomotor mappings and the generalization of this learning to intermediate starting locations demonstrated an interpolation of the two learned maps. This interpolation was a weighted average of the two learned visuomotor mappings, with the weighting sigmoidally dependent on starting location—a prediction made by a computational model of modular learning known as the "mixture of experts" [1]. These results provide evidence that the brain may employ a modular decomposition strategy during learning.

A general strategy for learning is to divide a complex task into simpler subtasks and learn each subtask with a separate module. This strategy has recently been formalized into a computational model of learning known as the mixture of experts [1], in which a set of expert modules each learn one of the subtasks and a gating module weights the contribution of each expert module's output to the final system output. The gating module bases its weighting of each expert on its estimate of the probability that this expert is the appropriate one to use for the current task. During learning, the gating module simultaneously learns to partition the task into subtasks while the expert modules learn these subtasks. Such modular decomposition has been proposed both as a model of high-level vision [4] and of the role of the basal ganglia during sensorimotor learning [5]. The mixture of experts model makes specific predictions regarding the nature of learning which have not been tested empirically. Here we test the hypothesis that the visuomotor system exhibits such modular decomposition during learning.

Previous studies have shown that the motor system is able to adapt to multiple different perturbations. Subjects adapt increasingly readily when repeatedly presented with two different prismatic displacements separated temporally [6, 7], a process which is mediated by posterior parietal cortex [8]. Similarly, subjects adapt to multiple perturbations if cued by gaze direction [9, 10, 11], body orientation [12], arm configuration [13], an auditory tone [14] or the feel of prism goggles [15, 16, 17]. One hypothesis to account for these studies is that multiple visuomotor mappings are stored simultaneously, suggesting a modular system. However, alternative explanations, such as a general increase in adaptability, or a single, non-modular system which is responsive to inputs from many modalities, cannot be ruled out from these studies. In particular it is not clear if the outputs of separate modules can be appropriately combined for contexts not already learned. Here we probe the existence of multiple modules by testing the specific predictions of a computational model of modular learning.

We investigated a learning paradigm in which the visual feedback of the hand during pointing movements was perturbed so that a single location in vi-

*Correspondence should be addressed to Zoubin Ghahramani, Department of Computer Science, University of Toronto, Canada M5G 3H5. Email: zoubin@cs.toronto.edu.

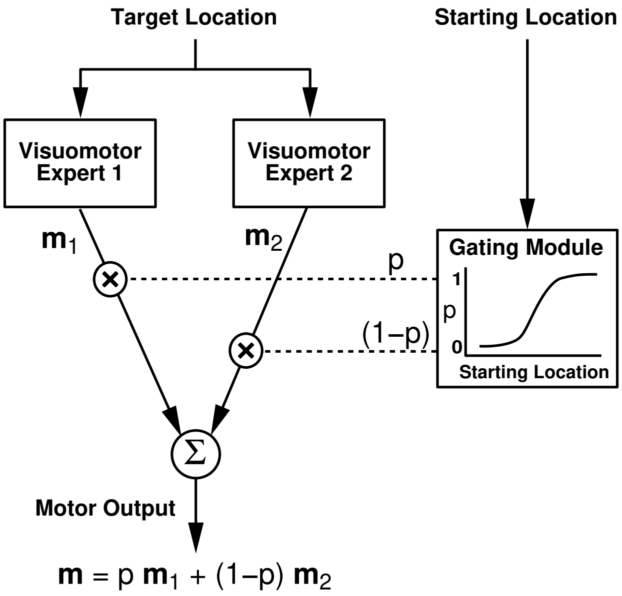

Figure 1. A modular decomposition model of visuomotor learning in which two different maps can be learned for the same visual target location. This represents the simplest instantiation of the hierarchical mixture of experts [2], having only one level and two experts. The model maps target and starting locations to motor outputs, $\mathbf{m}$, which could represent, for example, the final hand location or movement vector. Each expert learns a different mapping between target locations and motor outputs. The contribution of each expert's output, $\mathbf{m}_1$ and $\mathbf{m}_2$, to the final motor output, $\mathbf{m}$, is determined by the gating module's output, p. The output p reflects the probability that expert 1 is the correct module to use for a particular starting location—at p values of 1 or 0 the final output is determined solely by the output of expert 1 or expert 2 respectively, whereas at intermediate values of p both experts contribute to the final output. The logistic form of the gating module's output as a function of starting location can be derived by assuming that each expert learns the visuomotor map at one of the two starting locations—its preferred starting location—and that each expert is responsible for an equal size Gaussian region around this preferred starting location.

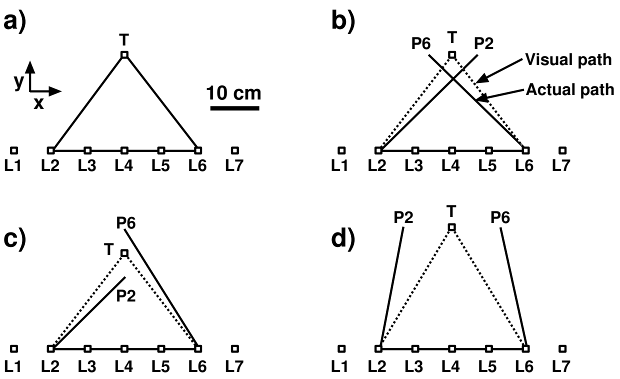

Figure 2. A schematic of the perturbations used for the four groups of subjects (a–d). Movements were made in the horizontal plane and the schematic shows the seven possible starting locations (L1–L7) and the target (T) seen from above. The solid lines indicate the actual path taken by the hand during the exposure phase; the dotted lines indicate the visually displayed path of the hand. For the control group (a) the two lines coincide everywhere as there was no perturbation and therefore no discrepancy between the visually displayed and actual hand location. For the perturbation groups a discrepancy between displayed and actual hand position was introduced (see Methods for details). The discrepancy was chosen so that subjects, in order to visually perceive their hand on target T, had to point to two different locations, P2 & P6, depending on whether the movement started from L2 or L6.

sual space was remapped to two different hand positions depending on the starting location of the movement (see Methods for details). This perturbation creates a conflict in the visuomotor map, the internal model [18] of the kinematics of the arm which captures the normally one-to-one relation between visually perceived and actual hand locations [19, 20, 21]. One way to resolve this conflict is to develop two separate visuomotor maps, the expert modules, each appropriate for one of the two starting locations (Figure 1). A separate mechanism, the gating module, then combines, based on the starting location of the movement, the outputs of the two visuomotor maps. The output of the gating module, which represents the weighting given to each visuomotor map for a given starting location, has a sigmoidal (logistic) shape, as a function of the starting location of the movement (Figure 1). This relationship results from the assumption that each expert is responsible for an equal variance Gaussian region around its preferred starting location [22], which corresponds to its receptive field. As in previous studies of the visuomotor system [23, 24, 25], the internal structure of the system can be probed by investigating the generalization properties in response to novel inputs, which in this case are the starting locations on which it has not been trained. The hallmark of a system with modular decomposition is the ability to learn both conflicting mappings, and to transition smoothly from one visuomotor map to the other in a sigmoidal fashion as the starting location is varied.

Subjects were exposed in a virtual reality setup to two different visuomotor perturbations, discrepancies between the actual and perceived hand location, depending from which of two possible starting locations the movement originated (L2 & L6 in Figure 2; see Methods). Although subjects were unaware of the perturbation, they showed significant adaptive changes in their pointing behaviour when starting

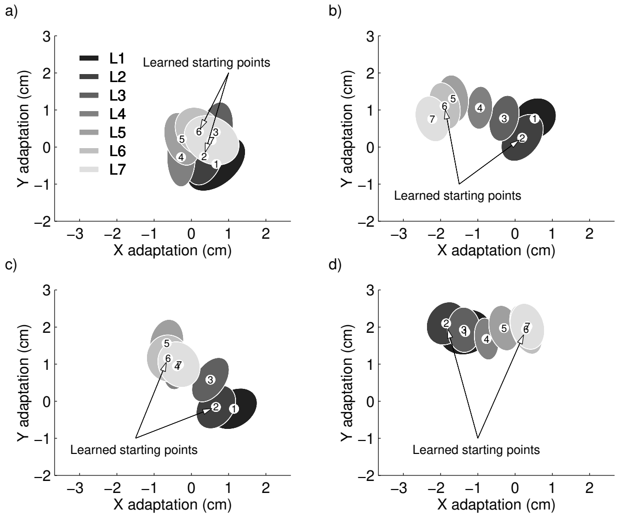

Figure 3. Adaptation of the visuomotor map for the control (a) and perturbation groups (b–d). For each starting location (L1–L7 denoted by shading) the 95% confidence ellipse for the change in pointing behaviour, induced by the visuomotor perturbation, is shown. For clarity, the centre of the ellipses, which represents the change in pointing behaviour, is also indicated numerically by the starting location (e.g. 3 corresponds to the change in pointing for movements starting from L3). The change in pointing corresponding to the learned starting points L2 & L6 are indicated by the arrows. For the perturbation groups, significant changes in pointing are seen, corresponding to partial adaptation to the remappings introduced. These changes in pointing smoothly shift as the starting location varied between L1 and L7. As well as the changes in response to the perturbation, there was a starting-point-independent movement overshoot for the perturbation groups, accounting for both the Y offset of the means in b and d and the additional right-to-left shift seen in c.

from locations L2 and L6 (Figure 3 b–d). The adaptation seen for movements from these two points was significantly different from each other ($P < 0.001$), showing that the subjects were able to learn two distinct remappings of the same point in visual space as a function of the starting location. Furthermore, as the starting location is varied between L1 and L7 a smooth transition can be seen in the change in pointing behaviour which reflects visuomotor learning (Figure 3 b–d).

We estimated the mixing proportion (p in Figure 1) by fitting the changes in pointing behaviour at each starting location to a weighted mixture of the adaptation observed for movements starting from L6 and L2. These estimates show a significant modulation over the starting locations (Figure 4) for groups b–d ($P < 0.001$), but as expected, not for the control group ($P > 0.05$) who showed no change in pointing behaviour. The modulation in groups b–d showed a significantly better fit to a logistic function, the mixing probabilities predicted by the modular decomposition model, than to a linear function ($P = 0.02$).

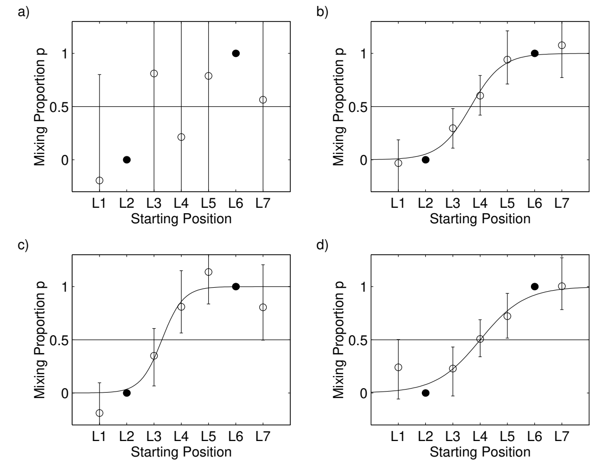

Figure 4. The mixing proportions with 95% confidence limits as a function of starting location for the (a) control and (b–d) perturbation groups. For the i^{th} starting location, the mixing proportion p is computed to minimize the distance between $\mathbf{v}_i$ and $p\mathbf{v}_6 + (1-p)\mathbf{v}_2$, where $\mathbf{v}_i$ is the mean adaptation vector for starting location Li. Using this criterion the values of p are fixed to be 0 and 1 at starting locations L2 and L6 (indicated by filled circles), respectively. The values of p at points other than L2 and L6 capture the form of the generalization as a function of the two learned mappings at L2 and L6. Confidence intervals were computed on this measure using bootstrap re-sampling [26]. The mixing proportions for the control group (a) did not differ significantly from the null hypothesis of equal mixing (indicated by the line at 0.5). For groups (b–d) a logistic function, $p(i) = 1/[1+\exp(a+bi)]$, representing the mixing probabilities predicted by the modular decomposition model (Figure 1), was fit to the mean mixing proportions (solid curve). All three fits were significant ($P < 0.001$) and the logistic function fit was a significantly better fit than linear regression over the ensemble data sets b–d ($F(15, 15) = 3.17$; $P = 0.02$).

The hypothesis of modular decomposition can be contrasted with models in which a single visuomotor transformation is computed. Models in which the transformation is based solely on the visual location of the target cannot account for the two mappings learned for the same point in visual space. Alternatively, a single visuomotor transformation may take in as inputs both the visual location of the target and the movement starting location. The manner in which such a single module would generalize to new starting locations depends crucially on the internal structure of the module. For example, a linear model [23] predicts a corresponding linear pattern of generalization, which was not observed in the data (Figure 4 b–d). Conversely, the internal structure of a single module may consist of units with receptive fields which are local both in visual space, as indicated by recent studies of generalization in the visuomotor map [25], and starting location. In this case, the predicted adaptation would be maximal at starting locations L2 and L6 and decay away from these points, a pattern also not supported by the data. Our study shows that two different maps can be learned for the same point in visual space and that the generalization to starting locations at which the subject was not exposed to the perturbation has the logistic relationship predicted by the mixture of experts model. These results provide evidence that modular decomposition is a feature of visuomotor learning.

These findings can be interpreted through the hypothesis that the visuomotor system maps visual vectors, pairs of target and starting locations, into movement vectors. Evidence for such vector-based coding has been obtained in neurophysiological studies which suggest that populations of cortical cells code for direction of movement [27, 28]. Similarly, it has been shown that a set of limb postures, which specify endpoints, can be achieved by stimulation of specific areas of the spinal cord, and that simultaneous stimulation of two such areas elicits a large repertoire of intermediate postures [29]. According to either of these hypotheses, our results show that learning two new visuomotor mappings, whether represented as vectors or postures, at the two starting locations, leads to a smooth sigmoidal generalization at intermediate locations. This generalization is consistent with a gradual mixing, modulated by starting location, of two separate neuronal populations, each of which has learned a different visuomotor mapping. This suggests a simple and plausible neural mechanism by which the modular learning observed could have arisen in the visuomotor system.

Methods

Thirty-two right-handed participants, who were naive to the purpose of the experiment and gave their informed consent, were randomly assigned to one of four groups: a, b, c and d.

Setup. Subjects sat at a large horizontal digitizing tablet with their head supported by a chin and forehead rest (a complete description of the setup can be found in [30]). The subjects held a digitizing mouse with their right index finger tip mounted on its cross hairs—direct view of their arm was prevented by a screen. The targets and feedback of hand position were presented as virtual images in the plane of the digitizing tablet, and therefore in the plane of the hand. This was achieved by projecting a computer display onto a horizontal rear projection screen suspended above the tablet. A horizontal front-reflecting semi-silvered mirror was placed face up midway between the screen and the tablet. The subjects viewed the reflected image of the rear projection screen by looking down at the mirror. By matching the screen-mirror distance to the mirror-tablet distance, all projected images appeared to be in the plane of the hand when viewed in the mirror. The targets were represented as 1 cm hollow squares and the hand position was displayed as a 6 mm filled white square, the cursor spot. The position of the hand was used on-line to update the position of this cursor spot at 60 Hz. The relation between the actual hand location and the hand cursor spot was computer controlled so as to allow arbitrary visuomotor perturbations. Therefore, the cursor spot could either accurately represent the true location of the hand or computer-controlled discrepancies between the cursor feedback and actual hand location could be introduced.

Paradigm. Subjects were asked to point to visually presented targets with their right hand. The experiment consisted of three parts: pre-exposure, exposure, and post-exposure. During pre- and post-exposure, subjects pointed to target T (10 repetitions for groups a–c; 15 repetitions for group d) in the absence of any visual feedback of the hand, starting from each of the 7 starting locations (L1–L7; Figure 2). This allowed the accuracy of pointing in the absence of visual feedback of hand location to be assessed for 7 starting locations.

During the exposure phase, subjects repeatedly traced out a visual triangle L2-L6-T-L6-L2-T-L2 forty times, thereby alternately pointing to the target from L2 and L6, while receiving feedback of hand location via the cursor spot. For the control group (Figure 2 a) the hand cursor spot accurately represented the actual hand position at all times. For the perturbation groups (Figure 2 b–d), displacements were surreptitiously introduced between the actual and visually displayed hand location. The displacement introduced increased linearly with distance from the starting location; the direction of the displacement varied between the groups. For movements made during the exposure phase the sign of the displacement was different for the two starting locations, L2 and L6. The dotted lines in (Figure 2 b–d) show the path taken by the visual feedback of the hand location and the solid lines the actual path taken by the hand. For example, for group b a discrepancy was introduced so that visual feedback of hand position was shifted to the left for movements made from L2, reaching a maximum discrepancy of 5 cm when the visual feedback of the hand was on target. However, for movements from L6 the visual feedback of hand position was shifted to the right, again reaching a maximum of 5 cm. Consequently, the single visual target location (T) was remapped to two distinct hand locations (P2 & P6; Figure 2) depending on whether the movement started from L2 or L6. Movements between L2 and L6 were unperturbed in all groups.

To assess learning and generalization to movements made from other starting locations, the subjects' change in pointing behaviour between the pre-exposure and post-exposure phases was analyzed for each starting location. For each subject and start location the average change in pointing position between the pre-exposure and post-exposure phases was calculated, along with the corresponding covariance matrices. The subjects' data were combined within each group for each starting location, obtaining the group mean change, along with the covariance matrix of the change, for each starting location. The change in pointing from each starting location was plotted as a 95% confidence ellipse centred on the mean change (Figure 3).

Acknowledgements

We thank M. I. Jordan for his support and G. E. Hinton, S. Goodbody and R. Lemon for comments on the manuscript. This project was supported by the Wellcome Trust, ATR Human Information Processing Research Laboratories, Siemens Corporation, the National Science Foundation, and the Office of Naval Research. Z. G. was supported by fellowships from the McDonnell-Pew Foundation and the Ontario Information Technology Research Centre.

References

[1] Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. Adaptive mixture of local experts. *Neural Computation* **3**, 79–87 (1991).

[2] Jordan, M. I. and Jacobs, R. A. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation* **6**, 181–214 (1994).

[3] Cacciatore, T. W. and Nowlan, S. J. Mixtures of controllers for jump linear and non-linear plants. In *Advances in Neural Information Processing Systems 6,* Cowan, J. D., Tesauro, G., and Alspector, J., editors, 719–726. Morgan Kaufmann Publishers, San Francisco, CA (1994).

[4] Jacobs, R. A., Jordan, M. I., and Barto, A. G. Task decomposition through competition in a modular connectionist architecture: The what and where vision tasks. *Cognitive Science* **15**(2), 219–250 (1991).

[5] Graybiel, A. M., Aosaki, T., Flaherty, A. W., and Kimura, M. The basal ganglia and adaptive motor control. *Science* **265**(5180), 1826–1831 (1994).

[6] McGonigle, B. and Flook, J. Long-term retention of single and multistate prismatic adaptation by humans. *Nature* **272**, 364–366 (1978).

[7] Welch, R. B., Bridgeman, B., Anand, S., and Browman, K. E. Alternating prism exposure causes dual adaptation and generalization to a novel displacement. *Perception and Psychophysics* **54**(2), 195–204 (1993).

[8] Clower, D. M., Hoffman, J. M., Votaw, J. R., Faber, T. L., Woods, R. P., and Alexander, G. E. Role of posterior parietal cortex in the recalibration of visually guided reaching. *Nature* **383**, 618–621 (1996).

[9] Kohler, I. Development and alterations of the perceptual world: conditioned sensations. *Proceedings of the Austrian Academy of Sciences* **227**, 1–118 (1951).

[10] Hay, J. C. and Pick, H. L. Gaze-contingent prism adaptation: Optical and motor factors. *J. of Experimental Psychology* **72**, 640–648 (1966).

[11] Shelhamer, M., Robinson, D. A., and Tan, H. S. Context-specific gain switching in the human vestibuloocular reflex. In *Annals of The New York Academy of Sciences,* Cohen, B., Tomko, D. L., and Guedry, F., editors, volume 656, 889–891. New York Academy of Sciences, New York (1991).

[12] Baker, J. F., Perlmutter, S. I., Peterson, B. W., Rude, S. A., and Robinson, F. R. Simultaneous opposing adaptive changes in cat vestibulo-ocular reflex directions for two body orientations. *Experimental Brain Reseach* **69**, 220–224 (1987).

[13] Gandolfo, F., Mussa-Ivaldi, F. A., and Bizzi, E. Motor learning by field approximation. *Proc. Nat. Acad. Sci. (USA)* **93**, 3843–3846 (1996).

[14] Kravitz, J. H. and Yaffe, F. Conditioned adaptation to prismatic displacement with a tone as the conditional stimulus. *Perception and Psychophysics* **12**, 305–308 (1972).

[15] Kravitz, J. H. Conditioned adaptation to prismatic displacement. *Perception and Psychophysics* **11**, 38–42 (1972).

[16] Welch, R. B. Discriminative conditioning of prism adaptation. *Perception and Psychophysics* **10**, 90–92 (1971).

[17] Martin, T. A., Keating, J. G., Goodkin, H. P., Bastian, A. J., and Thach, W. T. Throwing while looking through prisms. II. specificity and storeage of multiple gaze-throw calibrations. *Brain* **119**, 1199–1211 (1996).

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

[19] Andersen, R. A., Essick, C., and Siegel, R. Encoding of spatial location by posterior parietal neurons. *Science* **230**, 456–458 (1985).

[20] Soechting, J. F. and Flanders, M. Sensorimotor representations for pointing to targets in three-dimensional space. *J. Neurophysiology* **62**, 582–594 (1989).

[21] Kalaska, J. F. and Crammond, D. J. Cerebral cortical mechanisms of reaching movements. *Science* **255**, 1517–1523 (1992).

[22] Bridle, J. S. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In *Neurocomputing: algorithms, architectures, and applications,* Fougelman-Soulie, F. and Herault, J., editors. Springer-Verlag (1990).

[23] Bedford, F. Constraints on learning new mappings between perceptual dimensions. *J. of Experimental Psychology: Human Perception and Performance* **15**(2), 232–248 (1989).

[24] Imamizu, H., Uno, Y., and Kawato, M. Internal representations of the motor apparatus: Implications from generalization in visuomotor learning. *Journal of Experimental Psychology: Human Perception & Performance* **21**(5), 1174–1198 (1995).

[25] Ghahramani, Z., Wolpert, D. M., and Jordan, M. I. Generalization to local remappings of the visuomotor coordinate transformation. *J. Neuroscience* **16**(21), 7085–7096 (1996).

[26] Efron, B. *The Jacknife, the Bootstrap and Other Resampling Plans.* Society for Industrial and Applied Mathematics, Phildelphia, PA, (1982).

[27] Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986).

[28] Georgopoulos, A. P. Current issues in directional motor control. *Trends in Neurosci.* **18**(11), 506–510 (1995).

[29] Bizzi, E., Mussa-Ivaldi, F. A., and Giszter, S. Computations underlying the execution of movement: A biological perspective. *Science* **253**, 287–291 (1991).

[30] Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study . *Experimental Brain Research* **103**(3), 460–470 (1995).