

Interpolation and Extrapolation in Human Behavior and Neural Networks

Emmanuel Guigon

Abstract

■ Unlike most artificial systems, the brain is able to face situations that it has not learned or even encountered before. This ability is not in general echoed by the properties of most neural networks. Here, we show that neural computation based on least-square error learning between populations of intensity-coded neurons can explain interpolation and extrapolation

capacities of the nervous system in sensorimotor and cognitive tasks. We present simulations for function learning experiments, auditory–visual behavior, and visuomotor transformations. The results suggest that induction in human behavior, be it sensorimotor or cognitive, could arise from a common neural associative mechanism. ■

INTRODUCTION

There are many circumstances where biological organisms must be able to produce appropriate responses to stimuli beyond the range of their previous experience. This capacity to face novel situations using limited prior knowledge is called *induction*, and it forms the basis of many motor and cognitive skills (e.g., forecasting, decision making) that require one to discover a continuous relation between stimulus and response. The nature of induction processes has been studied in function learning experiments (DeLosh, Busemeyer, & McDaniel, 1997; Koh & Meyer, 1991; Carroll, 1963). In a typical paradigm, participants first learn a collection of associations between stimuli and responses drawn from arbitrary dimensions and are then asked to generate responses to untrained stimuli. The characteristics of these responses provide a measure of interpolation (Koh & Meyer, 1991) and extrapolation (DeLosh et al., 1997; Wagenaar & Sagaria, 1975) capacities. Induction can also be measured in sensorimotor behavior through localized distortions of sensorimotor relationships (Vetter, Goodbody, & Wolpert, 1999; Shinn-Cunningham, Durlach, & Held, 1998; Ghahramani, Wolpert, & Jordan, 1996; Ghahramani & Wolpert, 1997; Schor, Gleason, Maxwell, & Lunn, 1993; Bedford, 1989, 1993a, 1993b).

There are two fundamental issues in the study of induction: (1) Does induction result from a rule- or an associative example-based mechanism (Busemeyer, Byun, DeLosh, & McDaniel, 1997; DeLosh et al., 1997; Koh & Meyer, 1991; Brehmer, 1974; Carroll, 1963)? The rule-based approach was long favored with the development of a series of statistical models that accounted for

observations in probabilistic inference tasks (polynomial hypothesis testing, Brehmer, 1974; Carroll, 1963; log-polynomial adaptive regression, Koh & Meyer, 1991). The example-based approach was generally found to be inappropriate for extrapolation (DeLosh et al., 1997; Koh & Meyer, 1991; Carroll, 1963). Recently, DeLosh et al. (1997) showed that neither approach was sufficient on its own to explain extrapolation behavior in function learning, and so they proposed a hybrid model combining the two approaches. However, it is unclear how the nervous system can handle rule-based mechanisms. (2) Are there any similarities between the mechanisms involved in function learning and those involved in adaptation to sensorimotor distortions? It has been recognized that, in both cases, the adapted behavior is not a mere collection of input/output associations, but a true relation between entire dimensions of stimuli (DeLosh et al., 1997; Koh & Meyer, 1991; Bedford, 1989; Carroll, 1963). However, this observation is not sufficient to conclude that common neural operations exist.

In this article, we show that both sensorimotor and cognitive induction could result from an associative example-based mechanism based on computation with monotonically tuned neuronal populations (Guigon & Baraduc, 2002; Guigon, 2003). This single mechanism can explain interpolation and extrapolation capacities of humans in function learning (Busemeyer et al., 1997; DeLosh et al., 1997; Koh & Meyer, 1991), auditory–visual alignment (Shinn-Cunningham et al., 1998), and visuomotor behavior (Ghahramani & Wolpert, 1997). The present results elaborate previous modeling of proprioceptive–visual alignment (Guigon & Baraduc, 2002).

The principle of our model is to consider induction as a problem of how a neurally distributed representation of a function can be learned, and so the same general

architecture, in fact, a single-layer linear neural network, can be used to model induction in the sensorimotor and cognitive domain. We will first describe the general model and then show how it applies to three induction problems. In each case, we will give a brief overview of experimental methods and observations and then show that the model reproduces the main characteristics of these observations (shape of the induced function, time course of learning, response variability).

RESULTS

Function Learning

Interpolation

Interpolation capacities of the model were tested on the paradigm of Koh and Meyer (1991). In this study, participants were trained to discover functional relations between dimensions based on a restricted number of examples (practice pairs). Transfer pairs, which were in the range of the practice pairs, were used to assess the capacity for generalization. The participants learned three functions that related stimulus length (L , mm) and response duration (D , msec): (1) $D = 257.24L^{0.33}$ (power), (2) $D = 75 + 223.5 \ln L$ (logarithmic), and (3) $D = 453.5 + 10.9L$ (linear). Practice and transfer stimuli were [2.526, 4.492, 6.401, 8.030, 41.889, 52.447, 62.723, 75.019] and [13.084, 18.250, 23.416, 32.625], respectively. These functions, together with practice and transfer stimuli, are shown as insets in Figure 1.

Based on the preference of the subjects for power functions, Koh and Meyer (1991) suggested that sensory encoding resulted in a mapping in log–log coordinates. Thus, the network was trained to discover a relation between $\ln(L)$ and $\ln(D)$, and consequently, the power function is, in fact, linear. The dimension range was $[-1:54, 6:14]$ for $\ln(L)$ and defined by the current function for $\ln(D)$.

During the early phases of training, the model produced a linear mapping (Figure 1, left), related to initial pretraining on the identity mapping. After further training, interpolation was near perfect (Figure 1, right). Interestingly, the variable error (measured as the variability of the output under zero-mean additive Gaussian noise) was the same between practice and transfer pairs (Figure 2A). Furthermore, the variable errors for all three functions decreased during the course of training and were larger for the logarithmic function than for the other two (Figure 2B–D). Similar results were obtained in human subjects (Koh & Meyer, 1991).

Extrapolation

Extrapolation capacities of the model were tested on the paradigm of DeLosh et al. (1997). In this study, participants were trained to convert a dimension D into a dimension R for three functional relations: $R = 2.2D +$

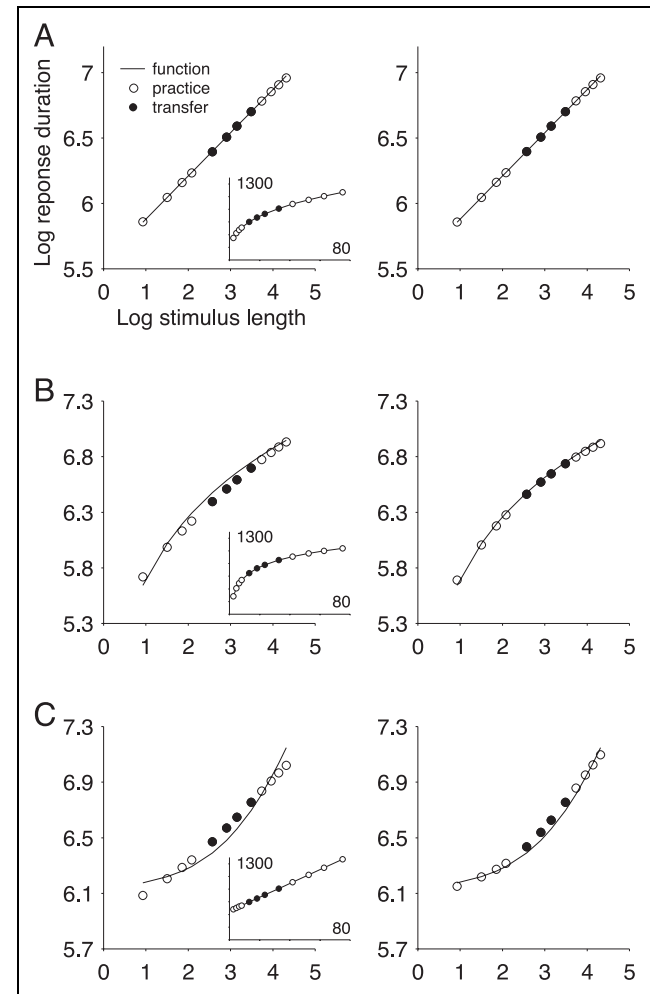


Figure 1. Interpolation in function learning: 500 trials (left column), 2500 trials (right column) ($\eta = 0.005$). (A) Power function. (B) Logarithmic function. (C) Linear function. Insets are the raw functions (length vs. duration).

30 (linear), $R = 200(1 - \exp(-D / 25))$ (exponential), and $R = 210 - (D - 50)^2 / 12$ (quadratic). Practice and transfer stimuli were [30.5, 36.0, 41.0, 46.5, 53.5, 59.0, 64.0, 69.5] and [1.0, 3.0, 5.0, 7.0, 9.0, 11.0, 13.0, 15.0, 17.0, 19.0, 21.0, 23.0, 25.0, 27.0, 29.0, 32.5, 35.0, 37.5, 40.0, 42.5, 45.0, 47.5, 50.0, 52.5, 55.0, 57.5, 60.0, 62.5, 65.0, 67.5, 71.0, 73.0, 75.0, 77.0, 79.0, 81.0, 83.0, 85.0, 87.0, 89.0, 91.0, 93.0, 95.0, 97.0, 99.0], respectively. The transfer pairs were outside the range of the practice pairs (Figure 3). The dimension range was $[-30, 130]$ for D and $[-80, 330]$ for R .

Linear, exponential, and quadratic functions were accurately approximated in the training range (Figure 3). Extrapolation outside this range was consistent with the shape of the target functions, although deviations were observed (Figure 3): underestimation of the linear function and overestimation of the exponential and quadratic functions. The linear function was learned faster than the exponential function, which in

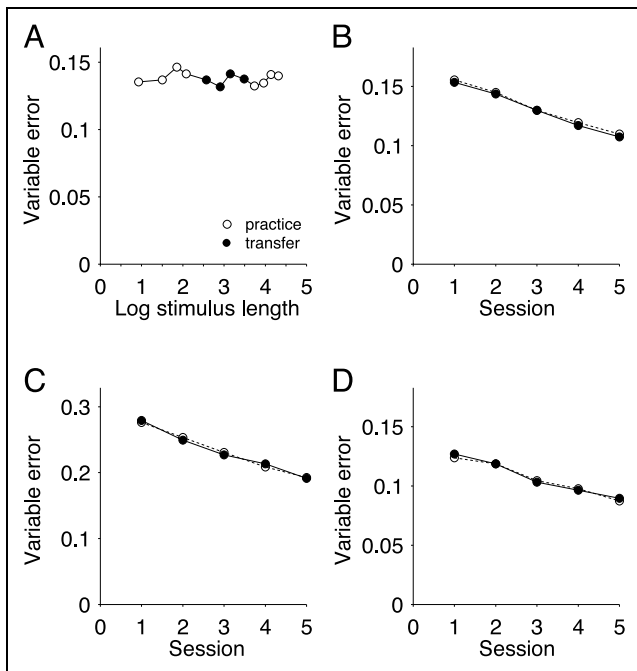


Figure 2. (A) Variable error in the model for the experiments of Koh and Meyer (power). Parameters are $\eta = 0.005$, additive Gaussian noise on the inputs (mean = 0, $SD = 0.1$, 500 repetitions). (B) Power function. (C) Logarithmic function. (D) Linear function. 1 session = 2000 trials.

turn was learned faster than the quadratic function (Figure 4). These results are consistent with experimental observations (DeLosh et al., 1997).

DeLosh et al. (1997) also showed that the density of stimuli within the training range had no significant influence on performance, a property found in our model as well. Results with low (8 stimuli), medium (20), and high (50) density are shown in Figure 4. The time course of error reduction was similar across the conditions.

Bussemeyer et al. (1997) reported further experiments in function learning. One of these is a comparison of acquisition of five monotonically increasing functions (Figure 5A). The order of learning difficulty was $LN > LG > PP > NP > PL$ (see figure legend). Order (Figure 5B,C) and acquisition errors in the model (Figure 5D,E) are compatible with experimental observations.

Induction in Auditory–Visual Behavior

In the sensorimotor domain, induction can be measured using localized distortions in natural sensory or sensory-to-motor congruences (e.g., Bedford, 1989). In a recent study, Shinn-Cunningham et al. (1998) trained subjects to interpret auditory localization cues arising from locations that differed from their normal spatial ones. Azimuthal location of auditory cues was transformed, so

that the perceived spatial location θ of a cue corresponded to the actual location $f_n(\theta)$ with

$$f_n(\theta) = \frac{1}{2} \tan^{-1} \left[\frac{2n \sin(2\theta)}{1 - n^2 + (1 + n^2)\cos(2\theta)} \right],$$

where n is the slope of the transform. Thirteen locations, spaced every 10° in azimuth from -60° to $+60^\circ$, were used as training locations. Although the subjects were trained to use nonlinear transformations ($n \neq 1$), the mean adapted response remained proportional to the normal cue location. In general, the slope of the response approached the best-fit slope of the

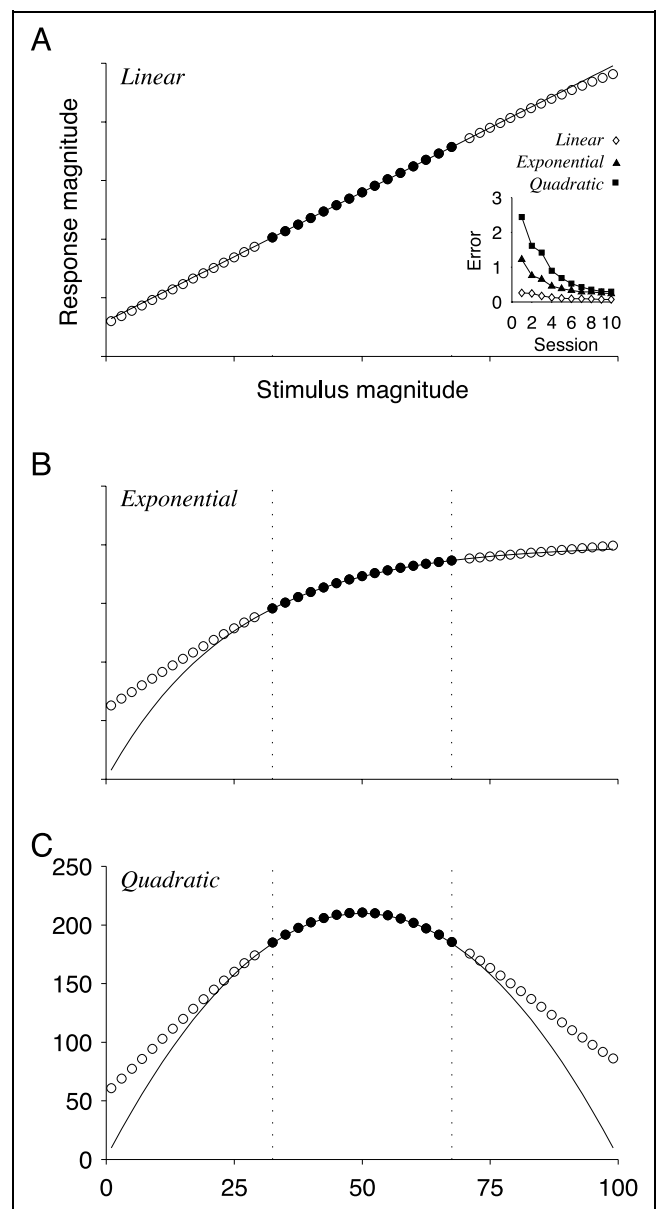


Figure 3. Extrapolation in function learning: 25,000 trials ($\eta = 0.04$). (A) Linear function. (B) Exponential function. (C) Quadratic function. Inset in the time of error during exposure. 1 session = 2000 trials.

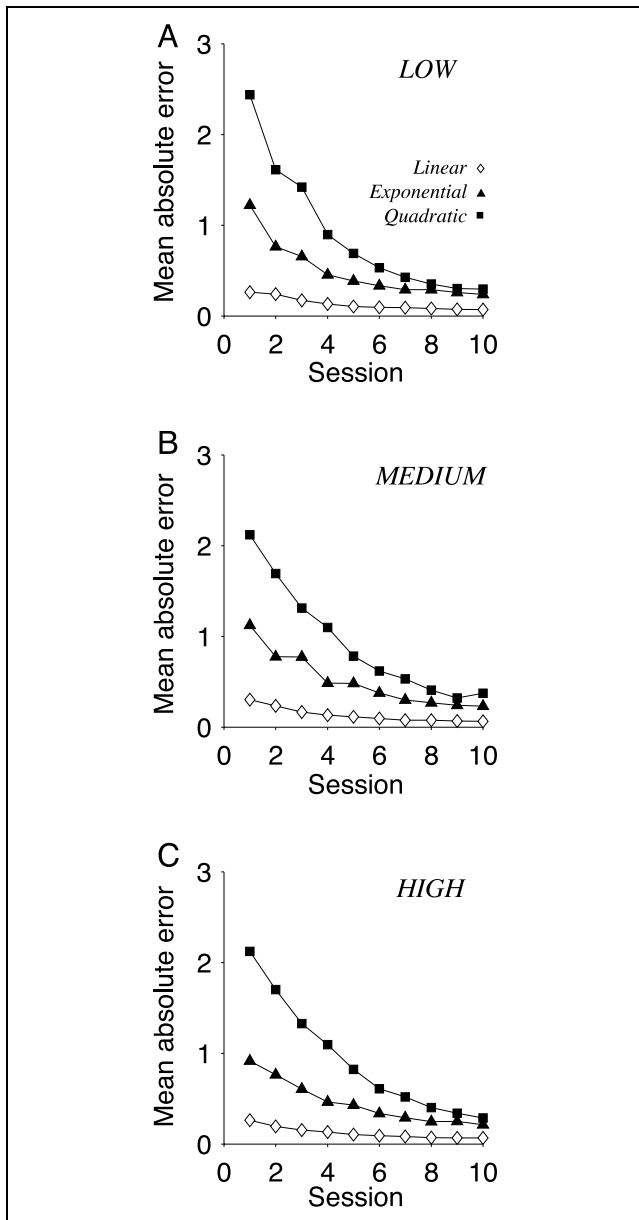


Figure 4. Influence of training stimulus density on the acquisition of linear, exponential, and quadratic functions: (A) low, (B) medium, and (C) high. 1 session = 2000 trials, $\eta = 0.001$.

training locations. The time course of the slope change was exponential.

The network learned the function $Y = f_n^{-1}(X)$ ($n = 3$), with training locations as described above, and the dimension range of $[-130^\circ, 130^\circ]$. Performance before learning is shown in Figure 6A. After 30 training blocks, the model response was linear and approximated the best-fit mapping defined by the training pairs (Figure 6B). After 1000 training blocks, the mapping became nonlinear (Figure 6C). More than 5000 training blocks were necessary to achieve the desired mapping (not shown). The time course of the slope change was exponential (Figure 6D).

Induction in Visuomotor Behavior

Ghahramani and Wolpert (1997) examined how induction occurred in a visual-to-motor transformation. They trained subjects to point to a target (T) from two starting positions, L2 and L6 (Figure 7, upper right). Three types of discrepancies were introduced between actual and visually perceived hand positions depending on the starting position (Figure 7, insets). For example, in the case of Figure 7D, the subjects should point more to the left of the target when starting from L2 and more to the right when starting from L6. Adaptation was measured as a change in pointing behavior from seven starting positions (L1–L7) following training. Ghahramani and Wolpert (1997) observed that the subjects actually adapted their pointing behavior, although they were

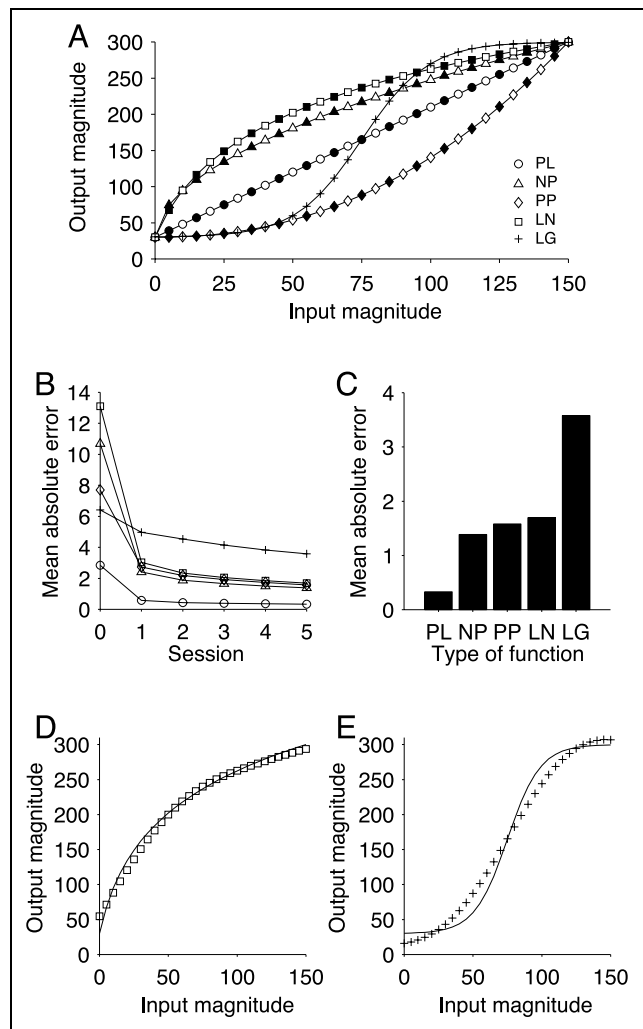


Figure 5. Acquisition of increasing functions. (A) PL = positive linear; NP = negatively accelerated power; PP = positively accelerated power; LN = logarithmic; LG = logistic. Open symbols are practice stimuli and closed symbols transfer stimuli. (B) Time course of error reduction. 1 session = 4000 trials. (C) Mean error after five sessions. (D) Errors for the LN function after five session. (E) Same as (D) for the LG function.

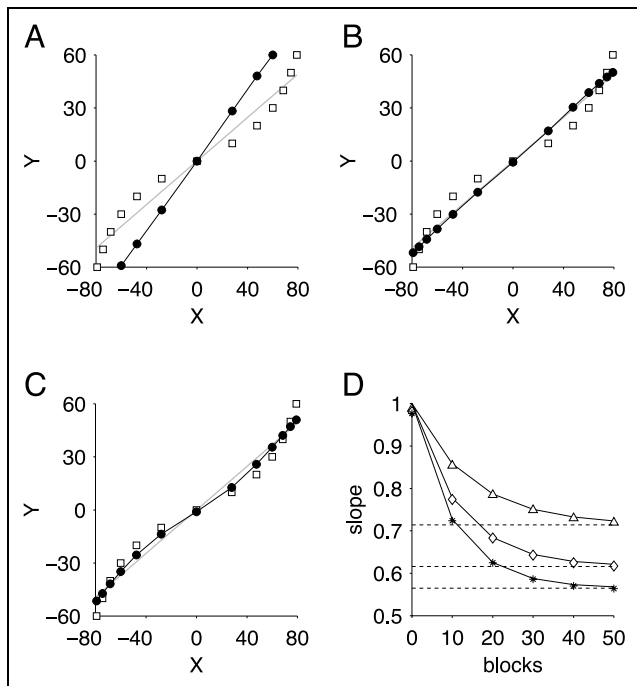


Figure 6. Adaptation to auditory-spatial distortions. Training pairs are depicted by \square (transformation slope $n = 3$) and adapted points by \bullet . In A, B, and C, the gray line is the best-fit line of the training pairs. (A) Pretraining performance. (B) Adaptation to the transformation after 30 training blocks. (C) Adaptation to the transformation after 1000 training blocks. (D) Time course of slope change (Δ : $n = 2$, \diamond : $n = 3$, $*$: $n = 4$). Symbols correspond to calculated points and lines to a monoexponential fit (time constants 14.6, 11.5, and 10.1 for $n = 2, 3$, and 4 , respectively). Dashed lines are the best-fit slope.

not aware of the perturbation. Adaptation occurred for the starting positions L2 and L6, but were smoothly generalized to fit the other starting positions.

In our model, pointing behavior was represented by the transformation $(P_i, T) \rightarrow P_f$, where P_i and P_f are the initial and final hand positions, respectively, and T the target position (see Discussion). The hand and target belong to a 2-D space. For the sake of simplicity, they were taken in the unit square $[0, 1] \times [0, 1]$. The monotonic model was used with P_i and T as inputs (cues) and P_f as output (target). The two dimensions of space were represented by distinct neural populations. We noted that actual movement direction could easily be obtained from P_i and P_f (e.g., using a single-layer linear network) and be used as a directional command for movement (Georgopoulos, 1995).

Pretraining consisted of 5000 random trials from the mapping $(P_i, T) \rightarrow T$. Exposure consisted of 1000 trials. Learning rate was $\eta = 0.01$. Adaptation is shown in the left column of Figure 7 and interpolation in the right column. Interpolation was defined at each position L1–L7 by a mixing proportion. At each position, L_i ($1 \leq i \leq 7$), the mixing proportion p_i is a weighted mixture of the adap-

tation observed for movements starting from L2 and L6 (Ghahramani & Wolpert, 1997) defined by $z_i = (1 - p_i)z_2 + p_i z_6$, where z_i is the adapted response at position L_i . The logistic shape of mixing proportions was found for $s = 0.025$, however, the relationship was linear for $s = 0.1$.

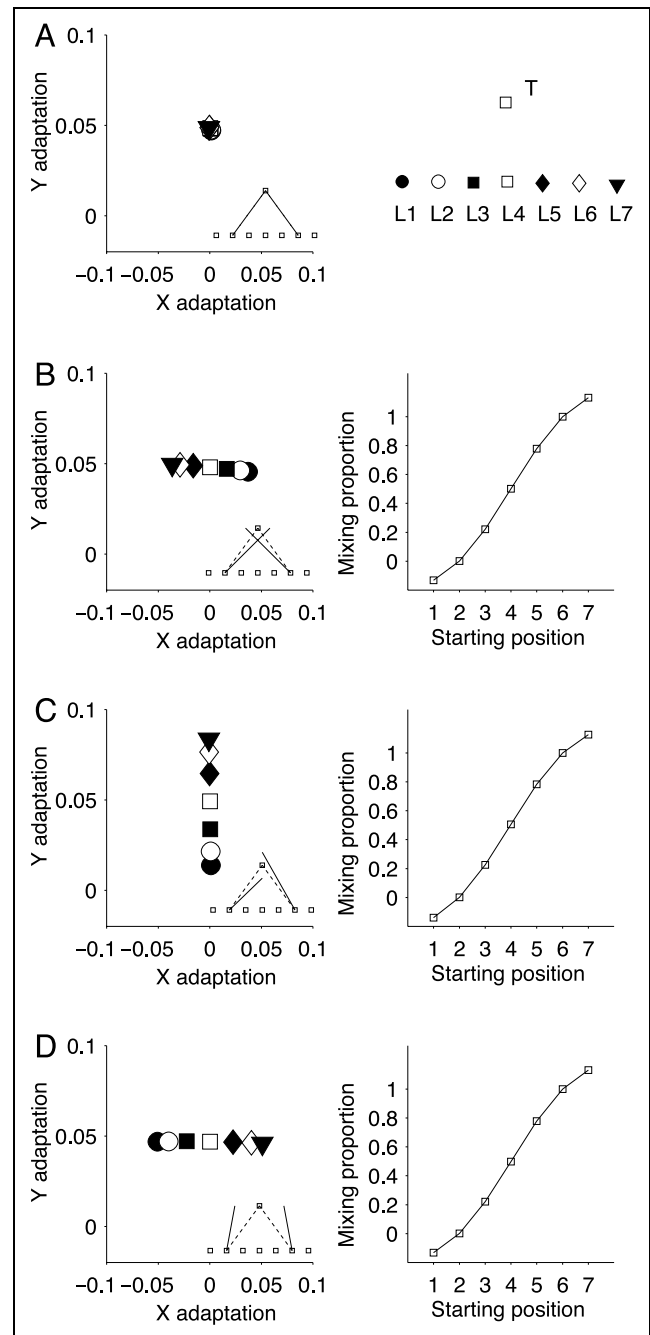


Figure 7. Experiment of Ghahramani and Wolpert (1997). Four pointing conditions are depicted (A, B, C, D, insets). Left column shows adaptation for the 7 starting positions (symbols). Right column shows mixing proportions. Positions were in the unit square $[0, 1] \times [0, 1]$. Starting positions were (0.3, 0.4), (0.37, 0.4), (0.43, 0.4), (0.5, 0.4), (0.56, 0.4), (0.63, 0.4), and (0.7, 0.4). Target was (0.5, 0.57). Target displacement was 0.05 in B and C, and 0.1 in D.

DISCUSSION

The model provides an economical approach to function learning experiments. It explains a wide range of results with a single associative mechanism and suggests that there is no need to invoke a rule to explain extrapolation capacities of human subjects. The failure of previous associative approaches is due to the use of the psychological similarity function borrowed from theories of category learning (Kruschke, 1992; Shepard, 1987; Nosofsky, 1986). This function quantifies the distance between stimulus and stored examples and is generally decreasing with distance. Thus, training effects remain localized within the width of the similarity function (Ghahramani et al., 1996), the variability in the experimentally measured responses being lower in the vicinity of the training regions than elsewhere, which is not supported by experimental observations (Ghahramani & Wolpert, 1997; Bedford, 1993a; Koh & Meyer, 1991). The monotonic response function, when considered as a pair of functions with positive and negative slopes, has no such drawback as it gives nonzero responses over the whole range of inputs and outputs. This property has a strong impact on learning capacities of neural networks, which has not been exploited before in the field of function learning. A single stimulus–response association can modify the whole synaptic structure of a network and induce well-structured responses to unexperienced stimuli.

Our model can only explain some of the basic properties of induction behavior, but it cannot account for more complex capacities. Humans can learn cyclic functions (Busemeyer et al., 1997), but the model cannot learn these functions. In fact, our model could learn cyclic functions by using a steeper response function (smaller s), or a network combining monotonic and broadly tuned inputs. The model cannot explain how either rule- or example-based representation can be combined (Erickson & Kruschke, 1998).

The models were designed to solve the problem of aligning two dimensions; however, in the case of visuomotor transformations (Ghahramani & Wolpert, 1997), this is problematic. A movement is defined by at least two parameters: an initial and a final position (or a direction and an amplitude). Thus, alignment should occur between these two parameters and an output dimension (e.g., actual final position or actual movement direction and amplitude). The nature of adaptation may depend on the choice of parameters. For our model, we chose positional parameters. An alternative representation of the transformation is $(P_i, T) \rightarrow (D, A)$, where D and A are the direction and amplitude of the movement, respectively. We have seen that movement direction (without amplitude information) can be readily obtained from the initial and final positions (by a single-layer linear mapping). However, neither the amplitude nor the vector $T - P_i$ can be easily obtained

in this way, which agrees with the notion that there is no separate representation of movement amplitude at the level of single neurons (e.g., Messier & Kalaska, 2000; Fu, Suarez, & Ebner, 1993). Thus, this alternative mapping may not be an appropriate representation of the visuomotor transformation. We cannot exclude the existence of other representations leading to similar predictions on the nature of adaptation.

A case in point is whether computation with populations of tuned neurons (i.e., with Gaussian tuning curves) could lead to similar observations. We have already shown that this is not generally the case for extrapolation (Guigon & Baraduc, 2002; Guigon, 2003), and so this cannot explain the results on function learning nor on auditory–visual alignment. In the case of visuomotor transformation, we observed that a Gaussian network does not properly compute positional transformation and the variability of responses (as measured following Gaussian noise injection) is more likely to be constant across positions (Ghahramani, Wolpert, & Jordan, 1995; Ghahramani & Wolpert, 1997), which is not compatible with a Gaussian model.

Ghahramani and Wolpert (1997) interpreted their results based on a sigmoid mixture of Gaussian experts. The idea was that there is a specialized network for each of the two training positions and each network learns the position-specific visuomotor relation with a Gaussian receptive field around the training position. The outputs of the two networks are then combined additively with proportions derived from a sigmoid representation of the initial starting position. Like in our model, Ghahramani and Wolpert found that generalization was based on the way arm posture is represented. However, they assumed that each expert network also uses a “localized” representation of posture, which has not been observed so far in the nervous system (see the discussion in Helms Tillery, Soechting, & Ebner, 1996). Alternatively, our model suggests that a single representation of posture can explain induction in visuomotor behavior.

The main parameter of this model is the slope s of the response function. Therefore, it is pertinent to determine what the value of s is for neurons in the central nervous system. We estimated s from the data of Squatrito and Maioli (1996) on eye position signals in the parietal cortex. To our knowledge, this is the only study that provides a quantitative description of variations in neuronal discharge with a postural parameter. The value of s was obtained by fitting a sigmoid on the average eye position signal (their Figure 13). We obtained values in the range 0.04–0.1 depending on the hypothesized maximal eccentricity (90–40°).

In summary, the present model suggests that sensorimotor and cognitive induction could result from a purely associative mechanism without the need to resort to additional computation (e.g., a linear extrapolation mechanism as described in DeLosh et al., 1997). This mechanism relies on built-in interpolation

and extrapolation capacities of populations of monotonically responding neurons (Guigon & Baraduc, 2002; Guigon, 2003) and provides an economical explanation to function learning experiments and sensorimotor adaptations.

METHODS

A neural representation of a real function $y = b(x)$ (x and $y \in [0, 1]$) was obtained as a linear mapping $\mathbf{y} = \mathbf{W}\mathbf{x}$, where \mathbf{W} is an $N \times 2N$ matrix, $\mathbf{x} \in \mathbb{R}^{2N}$ is the distributed representation of x , and $\mathbf{y} \in \mathbb{R}^N$ is an output vector that is a distributed representation of the scalar output y of the network (Guigon & Baraduc, 2002). We used $\mathbf{x} = \{x_i\}$ with

$$x_i = \frac{1}{1 + e^{\pm(x - \lambda_i)/s}},$$

where λ_i are recruitment thresholds (uniform in $[0, 1]$) and s the steepness of the response function. The vector \mathbf{x} had $2N$ components, N with s and N with $-s$. The purpose of the positive and negative populations is explained in Guigon and Baraduc (2002). The output vector \mathbf{y} was decoded, using

$$y = \frac{1}{N} \sum_{i=1}^N y_i,$$

which is an efficient estimator of the content of \mathbf{y} (see Guigon & Baraduc, 2002; Guigon, 2003). The weight matrix was identified by least-square error learning (rate η) for a given training set $\{x^t, y^t = b(x^t)\}$,

$$\Delta \mathbf{W}_{ij} = \eta x_j^t (y_i^{t*} - y_i^t),$$

where $\{x_j^t\}$ are the components of the input \mathbf{x}^t , that is, the distributed representation of x^t , $\{y_i^t\}$ are the components of the output $\mathbf{y}^t = \mathbf{W}\mathbf{x}^t$, $\{y_i^{t*}\}$ the components of the desired output \mathbf{y}^{t*} , the distributed representation of $b(x^t)$.

The same formalism applies, when $x \in [a, b]$, by defining a new variable $(x - a)/(b - a)$, new recruitment thresholds $\{(\lambda_i - a)/(b - a)\}$ and a new steepness $s/(b - a)$. The interval $[a, b]$ is called the dimension range, which is calculated in the following way. First, the task range is defined as the maximum range of the data. It is chosen to include the training stimuli, but otherwise, it is arbitrary. If the task range is $[0.2, 0.6]$ in normalized coordinates, then the dimension range is $[0, 1]$. The dimension range is larger than the maximal range of the inputs to avoid decoding biases near the border of the interval (Guigon & Baraduc, 2002; Guigon, 2003). In general, the results depend quantitatively, but not qualitatively, on the choice of the dimension range.

The network was first trained to reproduce the identity mapping $y = x$ corresponding to alignment of the dimensions. Pretraining consisted of 10,000 presentations of randomly chosen pairs of the identity mapping

($\eta = 0.01$). Then, the alignment was perturbed by assigning new outputs to a discrete ensemble of inputs (practice stimuli, randomized presentation). The resulting mapping was assessed on a set of transfer stimuli.

Two methodological points need to be discussed: (1) The model has a large number of parameters ($2N \times N$ weights) and a small number of training pairs. This overparameterization could lead to overfit the data. However, overfitting is restricted by the parameter s , which defines the slope of the sigmoid function. The larger s , the stronger the linear constraint on the learned function. (2) The number of presentations necessary to learn a mapping depends on the way least-square error learning is implemented and has no relation to the number of presentations required by humans. Faster gradient methods would lead to faster learning.

Parameters were in general $N = 50$, $s = 0.1$.

Acknowledgments

I thank Pierre Baraduc, Romain Brette, and Etienne Koechlin for fruitful discussions and Christine Lienhart for revising my English.

Reprint requests should be sent to Emmanuel Guigon, INSERM U483, Université Pierre et Marie Curie, Boîte 23, 9, quai Saint-Bernard, 75005 Paris, France, or via e-mail: guigon@ccr.jussieu.fr.

REFERENCES

- Bedford, F. (1989). Constraints on learning new mappings between perceptual dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, *15*, 232–248.
- Bedford, F. (1993a). Perceptual and cognitive spatial learning. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 517–530.
- Bedford, F. (1993b). Perceptual learning. In D. Medin (Ed.), *The psychology of learning and motivation* (vol. 30, pp. 1–60). New York: Academic Press.
- Brehmer, B. (1974). Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. *Organizational Behavior and Human Performance*, *11*, 1–27.
- Busemeyer, J., Byun, E., DeLosh, E., & McDaniel, M. (1997). Function learning based on experience with input–output pairs by humans and artificial neural networks. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts, and categories* (pp. 405–437). Hove: Psychology Press.
- Carroll, J. (1963). *Function learning: The learning of continuous functional mappings relating stimulus and response continua* (ETS RB 63–26). Princeton, NJ: Educational Testing Service.
- DeLosh, E., Busemeyer, J., & McDaniel, M. (1997). Extrapolation: The sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning Memory and Cognition*, *23*, 968–986.
- Erickson, M., & Kruschke, J. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.
- Fu, Q., Suarez, J., & Ebner, T. (1993). Neuronal specification of direction and distance during reaching movements in

- the superior precentral premotor area and primary motor cortex of monkeys. *Journal of Neurophysiology*, *70*, 2097–2116.
- Georgopoulos, A. (1995). Current issues in directional motor control. *Trends in Neuroscience*, *18*, 506–510.
- Ghahramani, Z., & Wolpert, D. (1997). Modular decomposition in visuomotor learning. *Nature*, *386*, 392–395.
- Ghahramani, Z., Wolpert, D., & Jordan, M. (1995). Computational structure of coordinate transformations: A generalization study. In D. Touretzky, G. Tesauro, & T. Leen (Eds.), *Advances in neural information processing systems* (vol. 7, pp. 1125–1132). Cambridge: MIT Press.
- Ghahramani, Z., Wolpert, D., & Jordan, M. (1996). Generalization to local remappings of the visuomotor coordinate transformation. *Journal of Neuroscience*, *16*, 7085–7096.
- Guigon, E. (2003). Computing with populations of monotonically tuned neurons. *Neural Computation*, *15*, 2115–2127.
- Guigon, E., & Baraduc, P. (2002). A neural model of perceptual–motor alignment. *Journal of Cognitive Neuroscience*, *14*, 538–549.
- Helms Tillery, S., Soechting, J., & Ebner, T. (1996). Somatosensory cortical activity in relation to arm posture: Nonuniform spatial tuning. *Journal of Neurophysiology*, *76*, 2423–2438.
- Koh, K., & Meyer, D. (1991). Function learning: Induction of continuous stimulus–response relations. *Journal of Experimental Psychology: Learning Memory and Cognition*, *17*, 811–836.
- Kruschke, J. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, *99*, 22–44.
- Messier, J., & Kalaska, J. (2000). Covariation of primate dorsal premotor cell activity with direction and amplitude during a memorized-delay reaching task. *Journal of Neurophysiology*, *84*, 152–165.
- Nosofsky, R. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–61.
- Schor, C., Gleason, G., Maxwell, J., & Lunn, R. (1993). Spatial aspects of vertical phoria adaptation. *Vision Research*, *33*, 73–84.
- Shepard, R. (1987). Evolution of a mesh between principles of the mind and regularities of the world. In J. Dupre (Ed.), *The latest on the best: Essays on evolution and optimality* (pp. 251–275). Cambridge: MIT Press.
- Shinn-Cunningham, B., Durlach, N., & Held, R. (1998). Adapting to supernormal auditory localization cues. II. Constraints on adaptation of mean response. *Journal of the Acoustical Society of America*, *103*, 3667–3676.
- Squatrito, S., & Maioli, M. (1996). Gaze field properties of eye position neurones in areas MST and 7a of the macaque monkey. *Visual Neuroscience*, *13*, 385–398.
- Vetter, P., Goodbody, S., & Wolpert, D. (1999). Evidence for an eye-centered spherical representation of the visuomotor map. *Journal of Neurophysiology*, *81*, 935–939.
- Wagenaar, W., & Sagaria, S. (1975). Misperception of exponential growth. *Perception and Psychophysics*, *18*, 416–422.