

- more by adults than by children? *Visual Cognit.* 1, 253-274
- 47 Solso, R.L. and McCarthy, J.E. (1981) Prototype formation of faces: a case of pseudo-memory *Br. J. Psychol.* 72, 499-503
- 48 Poggio, T. and Edelman, S. (1990) A network that learns to recognize three-dimensional objects *Nature* 343, 263-266
- 49 Wallis, G. and Rolls, E.T. (1997) A model of invariant object recognition in the visual system *Prog. Neurobiol.* 51, 167-194
- 50 Edelman, S. and Weinhall, D. (1991) A self-organising multiple-view representation of 3D objects *Biol. Cybern.* 64, 209-219
- 51 Foldiak, P. (1991) Learning invariance from transformation sequences *Neural Comput.* 3, 194-200
- 52 Wallis, G. and Baddeley, R. (1997) Optimal unsupervised learning in invariant object recognition *Neural Comput.* 9, 883-894
- 53 Miyashita, Y. (1993) Inferior temporal cortex: where visual perception meets memory *Annu. Rev. Neurosci.* 16, 245-263
- 54 Stryker, M.P. (1991) Temporal associations *Nature* 354, 108-109
- 55 Wallis, G. (1998) Spatio-temporal influences at the neural level of object recognition *Netw. Comput. Neural Syst.* 9, 265-278
- 56 Sinha, P. and Poggio, T. (1996) Role of learning in three-dimensional form perception *Nature* 384, 460-463
- 57 Wallis, G. (1998) Temporal order in object recognition learning *J. Biol. Syst.* 6, 299-313
- 58 Stone, J.V. (1998) Object recognition using spatio-temporal signatures *Vis. Res.* 38, 947-951
- 59 Johansson, G. (1973) Visual perception of biological motion and a model for its analysis *Percept. Psychophys.* 14, 201-211
- 60 Bülthoff, H.H. and Yuille, A. (1996) A Bayesian framework for the integration of visual modules, in *Attention and Performance Vol. XVI: Information Integration in Perception and Communication*, (McClelland, J. and Inui, T., eds), pp. 49-70, MIT Press
- 61 Gillner, S. and Mallot, H.A. (1998) Navigation and acquisition of spatial knowledge in a virtual maze *J. Cogn. Neurosci.* 10, 445-463

Computational studies of the development of functionally specialized neural modules

Robert A. Jacobs

Three hypotheses about the activity-dependent development of functionally specialized neural modules are discussed in this review. These hypotheses state that: (1) a combination of structure-function correspondences plus the use of competition between neural modules leads to functional specializations; (2) parcellation is due to a combination of neural selectionism, the idea that learning results from a stabilization of some neural connections and the elimination of others, and a locality constraint, which states that connections between nearby neurons are more easily stabilized than those between distant neurons; and (3) a temporal and spatial modulation of plasticity can induce higher functional development in later-developing parts of the nervous system relative to earlier-developing parts. All three hypotheses have been implemented and evaluated in computational models. Limitations of current neuroscientific methodologies mean that computer simulation provides one of the only tools available for evaluating and refining our large-scale theories of the development of functionally specialized neural modules.

It is often assumed in the brain and cognitive sciences that the brain is a collection of modules that are each specialized to perform a specific function. Occasionally this assumption is made explicit and used as the foundation of a theory about the brain's underlying structure^{1,2}. It is important to note, however, that even if this assumption provides an accurate characterization of adult human brains, recent evidence raises

questions about its accuracy with respect to newborns and young children^{3,4}. This evidence suggests that newborns' neocortices are less structurally differentiated than those of adults in the sense that neocortical regions are anatomically more similar in newborns than in adults. In addition, the evidence suggests that cognitive processes are less localized in newborns' nervous systems than those of adults. Consequently,

it seems sensible to conclude that some aspects of the modular organization of the adult brain arise developmentally. If so, then it is important to know whether the developmental processes that determine the functional properties of a neural module operate according to fixed genetic instructions or whether these processes are also experience sensitive.

The roles assigned to nature and nurture in the acquisition of functional specializations have been modified in recent years. An increasing number of investigators are hypothesizing that experience-dependent processes are more influential in determining a brain region's functional properties than was previously supposed⁴⁻¹⁰. Some of the supporting evidence comes from studies of cortical localizations of cognitive functions in human patients. For example, Neville¹¹ and her colleagues compared the behavior and event-related brain potentials (ERPs) of hearing adults and congenitally deaf adults during the performance of visual attentional tasks. They found that ERPs to foveal stimuli were similar in congenitally deaf and hearing adults; however, ERPs over superior temporal cortical areas to peripheral stimuli were two to three times larger in deaf than in hearing subjects. Deaf adults also responded faster than hearing subjects in tasks requiring detection of movement in peripheral stimuli, though response times did not differ when foveal stimuli were used. Neville hypothesized that the portion of the visual system that mediates the processing of peripheral stimuli may, through a process of competitive interactions, take over brain regions in the congenitally deaf that would normally be auditory cortical fields either in primary sensory or multimodal cortical areas.

Further evidence of the experience-dependent nature of the acquisition of functional specializations comes from the study of developmental neurobiology. For instance, Sur¹² and his colleagues provided visual inputs to the auditory system of ferrets by inducing retinal afferents to project to the medial geniculate nucleus (MGN), also referred to as auditory thalamus. Consequently, visually responsive cells were recorded in MGN. MGN projects to primary auditory cortex and visually responsive cells were also found in this region. These cells tended to have large receptive fields with roughly one-third of the fields being orientation-selective and a similar proportion being direction-selective. Similar to the fields of simple or complex cells in normal visual cortex, the oriented receptive fields had either separate or co-extensive ON and OFF zones. In addition, many cells were driven binocularly. According to O'Leary⁴, these and other results support the conclusion that 'primary sensory areas arise from regions of developing neocortex that are initially similar or to some extent pluripotent.'

This article reviews three hypotheses about the development of functionally specialized neural modules, as well as computational models that implement these hypotheses. The emphasis is on developmental mechanisms, especially on mechanisms that are largely or exclusively activity-dependent. There is much important research on developmental mechanisms controlling neural features such as cell fate, laminar and areal patterning, and axonal connectivity that are not activity-dependent. This research is not reviewed here (see, for example, Chenn *et al.*¹³). Additionally, the article emphasizes a broad range of mechanisms. Some

hypothesized mechanisms (e.g. the parcellation processes described below) have been studied more extensively than others in the neuroscientific and computational literatures. This article, however, attempts to give a relatively balanced presentation of three plausible hypotheses.

The first hypothesis is that a combination of structure-function correspondences plus the use of competition between neural modules leads to functional specializations¹⁴⁻¹⁶. This conjecture has been instantiated in a family of neural network architectures referred to as 'mixtures-of-experts' architectures^{8,17}. The second hypothesis combines the idea of neural selectionism, which speculates that learning results from a stabilization of some neural connections and the elimination of other connections⁵, and a locality constraint, which states that connections between nearby neurons are more easily stabilized than those between distant neurons^{18,19}. The end result of a selectionist process that is subject to a locality constraint is that nearby neurons tend to perform more similar functions than neurons that are far apart, and that nearby neurons tend to communicate whereas distant neurons are isolated from each other. Such a process was studied using a neural network model that strengthened some connections and pruned other connections during the course of learning²⁰. The final hypothesis is that a temporal and spatial modulation of plasticity can induce higher functional development in later-developing parts of the nervous system relative to earlier-developing parts⁷. A neural network with temporally and spatially modulated learning rates was used to evaluate this conjecture²¹.

Competition between neural modules

Several researchers have speculated that the development of functionally specialized neural modules is based on a combination of structure-function correspondences plus the existence of competition between modules¹⁴⁻¹⁶. This theory contains two important notions. The first notion is that there are structure-function correspondences in the brain. Because different brain regions have different structural properties (e.g. different patterns of connectivity among their neurons), different regions are best at performing different types of functions. The second notion is that brain regions compete for the ability to perform a set of tasks. Regions become functionally specialized due to the competition; that is, different regions learn to perform different functions. This competition, however, is biased by the structure-function correspondences; each region tends to win the competition for those functions for which its structure makes it particularly well suited.

My colleagues and I have developed a neural network architecture, known as a mixtures-of-experts (ME) architecture, that implements this hypothesis. The architecture, which is illustrated in Fig. 1, consists of two types of networks: expert networks and a gating network. The expert networks compete to learn the training patterns, where a training pattern consists of an input along with a desired output; the gating network mediates this competition. Whereas the expert networks have an arbitrary connectivity, the gating network is restricted to have as many output units as there are expert networks, and the activations of these output units must be nonnegative and sum to one.

The output of the entire architecture, denoted y , is the linear combination of the experts' outputs:

$$y = \sum_{i=1}^n g_i y_i \quad [1]$$

where y_i denotes the output of the i th expert network and g_i is the gating network output corresponding to the i th expert.

The learning process of the ME architecture combines aspects of competitive and associative learning. Mathematically, the architecture can be characterized as a probability model known as a conditional mixture density model. Mathematical descriptions can be found in Jacobs *et al.*¹⁷ Jordan and Jacobs²² and Peng, Jacobs and Tanner²³. Here I present an intuitive description. The connection strengths of the expert and gating networks are adjusted simultaneously during training. Each expert network's output is compared with the target output at each time step. The expert whose output most closely matches the target is called the winner of the competition; the other experts are called losers. An expert receives an amount of training information that is proportional to its relative performance on the training pattern. The winning expert receives a lot of information, and thus learns a lot about the current training pattern, whereas the losing experts receive little or no information, and thus learn little about the current pattern. The gating network receives information about the relative performances of the experts on the current pattern. It adjusts its connection strengths so that when the current input (or a similar input) recurs in the future the activation of its output unit corresponding to the winning expert will be larger (closer to one) and the activations of its remaining output units will be smaller (closer to zero).

The learning process has a positive feedback effect that forces different expert networks to learn different tasks. This effect relies on the fact that, in general, training patterns from the same task share a common underlying structure, whereas patterns from different tasks have different underlying structures. Suppose that at some instant in time, an expert has won the competition to learn some of the training patterns from one particular task. The expert will, therefore, have at least partial 'knowledge' of the structure of the task. Consequently, in the future it will be likely to win the competition for the remaining patterns from that task. The expert will thereby become specialized for performing the task. However, as a result of this specialization, this expert will be likely to perform poorly on patterns from other tasks - unless some tasks happen to be very similar. Thus other experts will be likely to win the competition for the patterns from other tasks. In this way, different experts win the competition to learn patterns from different tasks, and the experts become specialized for performing different tasks.

The ME architecture is sensitive to structure-function correspondences. Jacobs and Kosslyn²⁴ for example, considered the hypothesis that different subsystems of the brain are responsible for making categorical and coordinate visual judgments. Categorical judgments include classifying the identify of a stimulus (e.g. object A is a dog); coordinate judgments include identifying a visual stimulus as a particular exemplar (e.g. object A is Fido). Much of the information

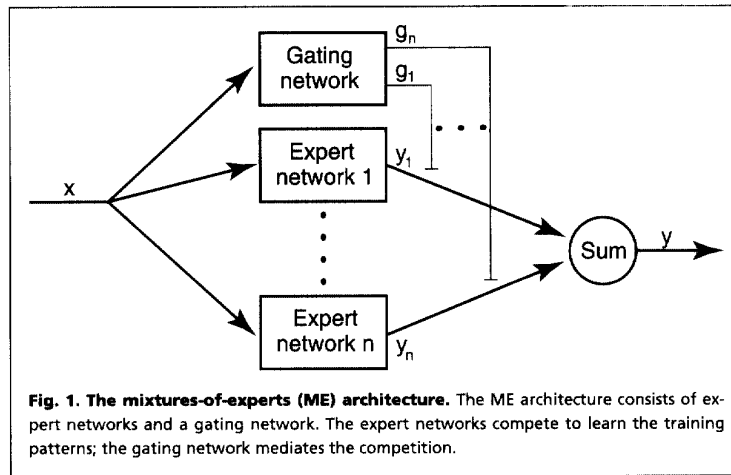


Fig. 1. The mixtures-of-experts (ME) architecture. The ME architecture consists of expert networks and a gating network. The expert networks compete to learn the training patterns; the gating network mediates the competition.

needed to make categorical judgments is irrelevant for making coordinate judgments, and vice versa. Categorization, for example, requires that various exemplars be grouped and treated as equivalent, whereas the identification of individual exemplars requires treating the instances as distinct. From an information processing viewpoint, it is logical that the brain might use different subsystems to make the two types of visual judgments. Jacobs and Kosslyn²⁴ Marsolek²⁵ and Marsolek, Schacter and Nicholas²⁶ reviewed experimental evidence from normal human subjects for a double dissociation between categorical and coordinate judgment tasks. Laeng found the same double dissociation for categorical and coordinate spatial relations judgments in a study using unilateral stroke patients²⁷.

Kosslyn *et al.*²⁸ speculated that there might be a structure-function relationship between receptive field sizes of neurons and visual judgments. Systems that make categorical visual judgments should be more efficient if they monitor visual neurons with small, non-overlapping receptive fields (it is important to note that populations of such neurons provide relatively low spatial resolution representations of visual images; see Fig. 2 for an intuitive explanation), whereas systems that make coordinate visual judgments

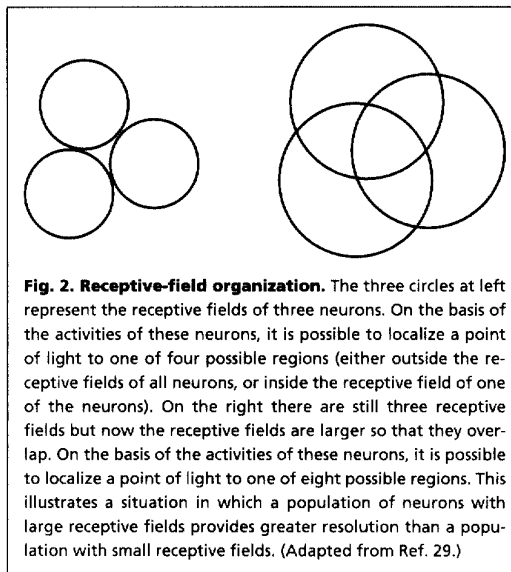


Fig. 2. Receptive-field organization. The three circles at left represent the receptive fields of three neurons. On the basis of the activities of these neurons, it is possible to localize a point of light to one of four possible regions (either outside the receptive fields of all neurons, or inside the receptive field of one of the neurons). On the right there are still three receptive fields but now the receptive fields are larger so that they overlap. On the basis of the activities of these neurons, it is possible to localize a point of light to one of eight possible regions. This illustrates a situation in which a population of neurons with large receptive fields provides greater resolution than a population with small receptive fields. (Adapted from Ref. 29.)

Box 1. ME architectures and inhibitory interactions

In a mixtures-of-experts (ME) architecture, the gating network plays a vital role in the sense that it determines the extent to which each expert network's output contributes to the output of the architecture as a whole. Note, however, that if one attempts to find a literal correspondence between the ME architecture and neural systems, there is no need to speculate about which specific neural structure might correspond to the gating network. This is because the ME architecture is equivalent to another architecture that contains expert networks but does not contain a gating network (that is, the two systems are exact notational variants of each other). The architecture without the gating network contains inhibitory connections among the expert networks so that each expert can suppress the outputs of the other experts. The strengths of these inhibitory connections are context dependent because they depend on the value of the current input pattern. (In the neural network literature, units whose connection strengths depend on the current input pattern are known as sigma-pi units.) At the end of training, the expert that was the winner of the competition in the context of the current input, or closely similar inputs, strongly suppresses the outputs of the other experts; experts that were losers of the competition do not suppress, or only weakly suppress, the other experts' outputs. The architecture without the gating network, but with inhibitory connections among the expert networks, is notable because the results of some experiments can be interpreted as suggesting that neural modules may use inhibitory interactions of this sort.

For instance, Gazzaniga (Ref. a) presented different visual inputs simultaneously to each hemisphere of a split-brain patient (this is a patient who has had the corpus callosum severed; this structure normally carries signals between the two cerebral hemispheres). The patient centered his or her gaze on a fixation point while a word was briefly presented so that half of the letters fell to one side of the point and half the letters fell to the other side. For example, if the word was target, then 'tar' fell in the patient's left visual field and was processed by the right hemisphere, whereas 'get' appeared in the right visual field and was processed by the left hemisphere. The patient was then presented with four letter strings and asked to point to the one that matched the visual input. The patient consistently pointed to the string that was presented to the right visual field regardless of which hand the patient used to perform the task. Glass *et al.* (Ref. b) interpreted these results as indicating that when there is a conflict between two plausible responses in this task, the left hemisphere inhibits the outputs of the right hemisphere and assumes motor control of both the left and right hands in making the response.

References

- a Gazzaniga, M.S. (1977) Consistency and diversity in brain organization *Ann. New York Acad. Sci.* 299, 415-423
- b Glass, A.L., Holyoak, K.J. and Santa, J.L. (1979) *Cognition* Addison-Wesley

should be more efficient if they monitor neurons with large, overlapping receptive fields (populations of such neurons provide high resolution representations of visual images). Jacobs and Kosslyn²⁴ used computer simulations to evaluate the proposed structure-function relationship. In brief, we trained neural networks to identify each visual stimulus as a member of a particular category ('shape category task') or to identify a stimulus as a particular exemplar ('shape coordinate task'). Networks did not view the visual stimuli directly; the stimuli were filtered through Gaussian units with restricted receptive fields. It was found that, indeed, the category task was learned faster when stimuli were filtered through units whose receptive fields were relatively small, whereas the coordinate task was learned faster when the stimuli were filtered through units whose receptive fields were large. When using an ME architecture, expert networks with small receptive fields tended to win the competition for the category task, whereas the coordinate task tended to be won by experts with large receptive fields. This set of simulations supports the hypothesized set of functional specializations and structure-function relationships by showing that these specializations and relationships are computationally efficient.

If competition plays a role in the acquisition of functional specializations in biological systems in the way that is suggested by the mixtures-of-experts architecture, then at least two predictions follow. The first is that there must exist initial differences between competing neural modules. For example, modules may initially differ in the information carried by the inputs they receive. Modules are, at least in part, potentially capable of displaying a wide range of

functional properties. As evidenced in the studies of Neville¹¹ and Sur *et al.*¹² discussed above, the input to a module is a powerful bias that helps determine the nature of the processing that the module ultimately performs. Simulation results with the ME architecture suggest that other differences between neural modules may also be important. Differences in receptive field size, number of processing units, and connectivity among the processing units may all serve to bias the relative learning performances of a collection of modules, thereby biasing their functional specializations. A second prediction is that neural modules should enforce the outcome of a competition through a set of adaptable inhibitory interactions that allow modules to suppress the outputs of other modules. This prediction is discussed more fully in Box 1.

Parcellation processes

Some investigators have hypothesized that functionally specialized neural modules develop via parcellation, a developmental process that produces regressive events such as the elimination of synapses and axon collaterals^{30,31}. Data suggests that humans lose approximately half of their synapses in cortical areas during the first few years of life³². The parcellation conjecture states that in postnatal development parcellation results in the elimination of neural connections such that previously combined information processing pathways or structures becoming segregated into relatively isolated modules.

At least two factors might influence the selection of synapses that are maintained versus those that are eliminated. First, according to the trophic theory of neural connections,

Box 2. Parcellation and the development of columnar structure

Columnar structures in the mammalian visual cortex, such as ocular-dominance columns or orientation columns, are often hypothesized to develop via a parcellation process. Miller (Ref. a) provides a review of computational models of the development of columnar structures in primary visual cortex, along with a mathematical analysis of these models. Here, we briefly outline a model of the development of ocular dominance columns due to Miller, Keller and Stryker (Ref. b). This model contains features that are commonly also found in other models. The model contains three sets of units. One set corresponds to neurons in the lateral geniculate nucleus (LGN) whose receptive fields are in the right eye, one set corresponds to LGN neurons whose receptive fields are in the left eye, and the final set corresponds to neurons in primary visual cortex (V1). The model contains four sets of parameters: (1) An arbor function $A(x-\alpha)$ gives the number of synapses between an LGN unit at location α and a V1 unit at location x ; (2) A cortical interaction function $I(x-y)$ describes the lateral influences among V1 units at locations x and y . These influences may occur through synaptic connections or through the diffusion of modulatory substances; (3) The patterns of activity in the afferent LGN units. The correlation functions C^{LL} , C^{RR} , C^{LR} and C^{RL} describe correlations in activity between afferents serving the same eye (C^{LL} and C^{RR}) or different eyes (C^{LR} and C^{RL}); (4) A parameter, denoted decay, that limits the total synaptic strength supported by a V1 unit. This latter parameter is used in the Hebbian learning rule that adapts the synaptic strengths associated with the connections between LGN and V1 units. At a coarse level of description, this rule may be written:

$$\Delta s = (\text{post} \times \text{pre}) - \text{decay} \quad [2]$$

where Δs is the change in synaptic strength s during a short time period, and post and pre are measures of postsynaptic and presynaptic activities during this period. Hebbian learning rules of this sort are sensitive to the correlational structure of the afferents in the sense that connections from afferent units with correlated activities tend to develop correlated synaptic strengths, whereas connections from afferent units with negatively correlated activities tend to develop synaptic strengths that are negatively correlated.

Miller *et al.* have shown that, over a wide range of parameter values, initially binocular V1 units tend to develop monocular receptive fields, and that nearby units tend to have receptive fields in the same eye. That is, the model develops ocular dominance columns that resemble those found in primate V1. Moreover, the model can also be used to mimic experimental findings collected under conditions of monocular deprivation and findings suggesting a critical period for the development of ocular dominance columns.

References

- a Miller, K.D. (1995) Receptive fields and maps in the visual cortex: models of ocular dominance and orientation columns, in *Models of Neural Networks* (Vol. 3) (Domany, E., van Hemmen, J.L. and Schulten, K., eds), pp. 55-78, Springer-Verlag
- b Miller, K.D., Keller, J.B. and Stryker, M.P. (1989) Ocular dominance column development: analysis and simulation *Science* 245, 605-615

presynaptic neurons compete for a limited supply of molecules, such as nerve growth factor, that are released by postsynaptic neurons. Presynaptic terminals that acquire sufficient amounts of these molecules maintain their connection to the postsynaptic neuron, whereas those that do not successfully compete for the molecules tend to lose their connection³³. Experimental evidence indicates that this competitive mechanism is capable of modulating the size of axonal and dendritic arbors as well as the size of neuronal populations. The distribution of trophic factors may be seen as a mechanism for implementing neural selectionism, the idea that learning results from a stabilization of some neural connections and the elimination of other connections⁵. An interesting hypothesis is that the competition for trophic factors is Hebbian, meaning that the amount of trophic factor available to a presynaptic cell is increased by near synchronous activity of presynaptic and postsynaptic cells^{34,35}. Second, the selection of synapses that are maintained or eliminated may be biased by a locality constraint such that connections between nearby neurons are more likely to be maintained than those between distant neurons, all else being equal. Barlow¹⁸ and Cowey¹⁹ have extensively discussed the role a locality constraint might play in determining neural organization.

Jordan and I presented a computational model that exhibits parcellation through the use of a Hebb-like learning process that is biased by a locality constraint²⁰. The learning process adapted the strengths of the connections of neural networks so that the networks more closely produced the

target output pattern for each input pattern. In addition, the learning process was biased such that the strengths of connections between nearby units were more likely to grow large in magnitude than those of connections between distant units. Simulation results showed that at the end of training nearby units performed more similar functions than distant units, and that nearby units communicated with each other whereas distant units were isolated from each other. Overall, the model provides support for the conjecture that functionally specialized neural modules can result from a selectionist learning process which is subject to a locality constraint.

Although only given a small amount of attention here, parcellation processes are among the most studied activity-dependent developmental processes in the neuroscientific and computational literatures. One area in which the role of parcellation processes seems to be particularly well understood is the development of columnar structures in primary visual cortex (see Box 2).

An open question regarding parcellation processes concerns the possible functional differences between different types of processes. Consider a neural region that initially has full connectivity; that is, each neuron is initially connected to all other neurons. One possibility is that parcellation occurs as postulated by the trophic theory of neural connections described above. A developmental process based on Hebbian principles leads to the maintenance of some neural connections and the retraction of other connections. Alternatively, parcellation could occur via a Hebbian learning

Box 3. Less is more

The notion of a 'wave of plasticity' is closely related to a set of ideas that have received considerable attention recently. Turkewitz and Kenny (Ref. a) speculated that the limited perceptual abilities of newborns should not be regarded as deficiencies that newborns need to overcome during the course of development. Rather, these limitations should be thought of as part of an adaptive developmental stage that facilitates the organization of sensory systems and provides a basis for subsequent perceptual and cognitive development. Newport (Ref. b) proposed a theory, known as the 'less is more' hypothesis, that is consistent with this framework in order to account for empirical data showing that people learn a language better if they are exposed to the language early in life versus late in life. The 'less is more' hypothesis speculates that the limited memory and attentional abilities of children are adaptive in the sense that they make it easier for children to identify the constituent parts that form a language.

Elman (Ref. c) implemented the less-is-more hypothesis in a neural network that was trained to process sentences from an artificial language. This language included linguistic features such as relative clauses, number agreement, and several types of verb argument structure. The neural network included recurrent

connections that allowed a subset of units to act as a short-term or working memory. The working memory capacity of the network could be manipulated by corrupting the activities of these units with noise. Simulation results suggest that networks with large working memory capacity cannot learn the underlying structure of the language. In contrast, networks with a working memory capacity that is initially small but gradually increases during the course of training can successfully learn the language. Elman concluded that the computational model supports the less-is-more hypothesis because it suggests that rather than being a limitation, developmental restrictions on resources may constitute a necessary prerequisite for mastering certain complex domains.

References

- a Turkewitz, G. and Kenny, P.A. (1982) Limitations on input as a basis for neural organization and perceptual development: a preliminary theoretical statement *Dev. Psychobiol.* 15, 357-368
- b Newport, E.L. (1990) Maturation constraints on language learning *Cognit. Sci.* 14, 11-28
- c Elman, J.L. (1993) Learning and development in neural networks: the importance of starting small *Cognition* 48, 71-99

rule that strengthens some anatomically fixed neural connections and weakens other connections. Miller³⁶ has shown the equivalence between an axonal sprouting-and-retraction model and a Hebbian synaptic plasticity model under certain conditions in the sense that the two models both maximize a metric that is roughly the sum of the correlations between the activities of synapses on the same and neighboring cortical cells. Consequently, other metrics or more detailed models are needed for theoretical distinctions to be drawn between sprouting-and-retraction models and Hebbian synaptic strength modification models.

The discussion in this section and in the previous section both emphasize the role of neural competition. Neural competition has been hypothesized to take place at many different scales in biological nervous systems. For example, competition takes place on a large scale in the mixtures-of-experts architecture in which the units of the winning competing module are allowed to learn about a data item, whereas the units of other modules are prevented from learning about that item. In contrast, parcellation hypotheses rely on neural competition on a small scale; the stabilization or strengthening of some connections between presynaptic neurons and a postsynaptic neuron is frequently accompanied by the loss or weakening of other connections to the postsynaptic neuron. Whereas there is much neuroscientific evidence suggesting the existence of neural competition at a small scale, the evidence showing neural competition at a large scale is relatively sparse.

'Wave of plasticity' hypothesis

Learning theorists have speculated that a good way to achieve large-scale systems that can solve difficult problems is to endow these systems with a 'bootstrapping' ability; that is, the systems should be able to use the solutions to simpler tasks as 'building blocks' for more difficult tasks^{37,38}. The development of biological nervous systems is often charac-

terized as using a bootstrapping strategy. Greenough, Black and Wallace⁷ hypothesized that asynchrony in brain development serves the useful function of 'stage setting'. The developmental schedule for the maturation of different brain regions is staggered such that neural systems that develop relatively early provide a suitable framework for the development of later, experience-sensitive systems (see Box 3). Experimental evidence consistent with this position comes from Harwerth *et al.*³⁹ These investigators performed behavioral studies of sensitive periods for visual development in monkeys. Their results suggest that these sensitive periods are organized into a hierarchy in which the visual functions requiring information processing in the peripheral portions of the visual system have shorter sensitive periods than those requiring more central processing.

Shrager and Johnson²¹ used the term 'wave of plasticity' to refer to the phenomenon that the locus of maximum neural plasticity seems to begin in the primary sensory and motor areas and then moves over time towards the secondary and parietal association areas and finally to the frontal regions. They investigated the possible consequences of such a wave by using a modified version of a computational model developed by Kerszberg, Dehaene and Changeux⁴⁰. A distinctive feature of this model is that synapses compete for a limited amount of trophic factor, as described above. Synapses that receive less trophic factor atrophy over time, whereas synapses that receive more trophic factor eventually reach stability. A Hebbian association rule directs the transportation of trophic factor from interneural space to synapses. Thus the period of plasticity depends upon the transportation of trophic factor, and may differ from one part of the network to another in accord with the activity-dependent diffusion of this factor. A second distinctive feature of the model by Kerszberg *et al.* is that it is possible to characterize the performance of each unit as a logic function both before and after training of a network.

In this way, it is possible to evaluate the simplicity or complexity of each unit's processing.

The network simulated by Shrager and Johnson²¹ consisted of a two-dimensional array of units in which each unit initially received inputs from two afferent units and from other units in the array (see Fig. 3). In the simulation most relevant to our purposes, they manipulated the distribution of trophic factor so as to produce a wave of plasticity; the locus of maximum plasticity moved across the array from left to right over time. The simulation produced a number of interesting findings. First, more synapses were eliminated in networks that contained a wave of plasticity than in those that did not. Synapses in networks containing a wave of plasticity that were far from the locus of maximum plasticity tended to receive little trophic factor and, thus, were often eliminated. Second, there was a tendency toward more complex processing in later-developing parts of the array. Shrager and Johnson explained this as follows. Let the set LEFT denote units in earlier-developing parts of the network, and let the set NEXT denote units in later-developing parts. The units in LEFT, initially more plastic than those in NEXT, determined their functions on the basis of the activities of the two afferent units. Over time, as the wave of plasticity moved to the right, the units in LEFT became relatively fixed in their function, and the units in NEXT became more plastic. Note that the units in NEXT, in addition to receiving the input coming from the afferent units, also received the activities of the units in LEFT. Consequently, they tended to develop more complex functions than those in LEFT. Third, networks with slow waves of plasticity tended to show compact clusters of units with simple functions or with complex functions, whereas networks with fast waves of plasticity had a more widely dispersed collection of units with complex functions.

The 'wave of plasticity' hypothesis is difficult to study experimentally. There are relatively few neuroscientific investigations that address it directly. Nonetheless, because it seems reasonable to assume that complex perceptual and cognitive information processing develops on the basis of simpler information processing, it may be among the most promising theories concerning the development of higher cortical areas.

Conclusion

This article has reviewed three hypotheses about the activity-dependent development of functionally specialized neural modules. These hypotheses state that: (1) a combination of structure-function correspondences and competition between neural modules leads to functional specializations; (2) parcellation is due to a combination of neural selectionism and a locality constraint, which states that connections between nearby neurons are more easily stabilized than those between distant neurons; and (3) a temporal and spatial modulation of plasticity can induce higher functional development in later-developing parts of the nervous system relative to earlier-developing parts. These hypotheses are not mutually-exclusive; different hypotheses may prove useful for understanding different neural and cognitive phenomenon. All three hypotheses have been implemented and evaluated in computational models. An impor-

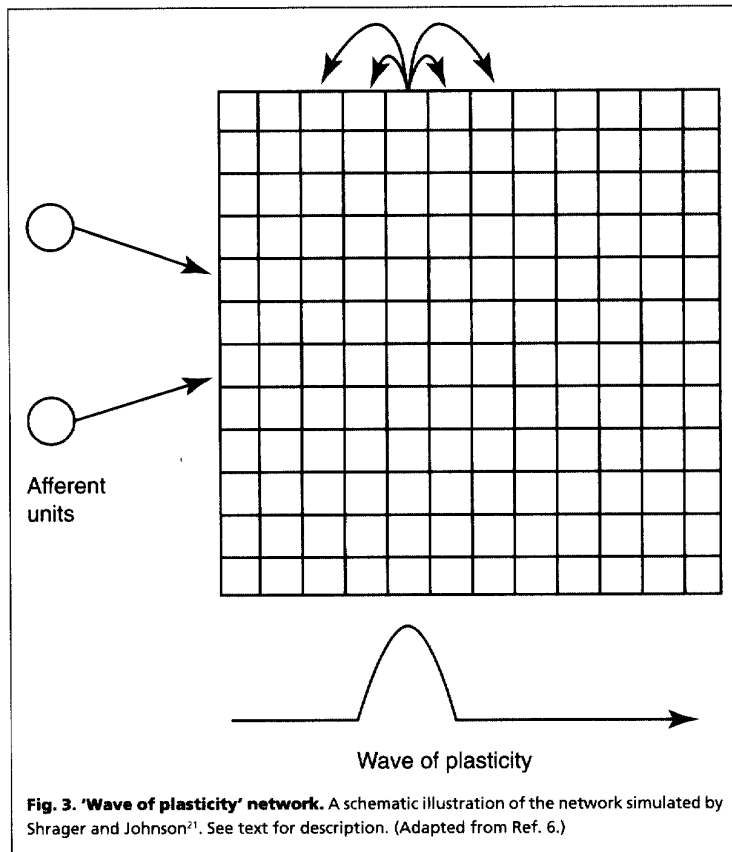


Fig. 3. 'Wave of plasticity' network. A schematic illustration of the network simulated by Shrager and Johnson²¹. See text for description. (Adapted from Ref. 6.)

tant role of these models is to specify the hypotheses in a more explicit manner and in greater detail than would otherwise be the case. In addition, the models are useful for evaluating the logic of the hypotheses, and for exploring possible elaborations of the hypotheses.

The hypotheses can be regarded as 'large scale' theories because they concern the changing characteristics of large portions of the brain over long periods of time. A drawback of such large-scale theories is that they often cannot be fully studied and evaluated using current neuroscientific methodologies. Most techniques for studying brain activity

Outstanding questions

- Neural competition has been conjectured to take place at a large scale (as in competition between neural modules) and at a small scale (as in competition between synapses of a postsynaptic neuron). What, if anything, do these forms of competition have in common?
- What are the functional differences between the situation in which there is an initial over-production of synaptic connections followed by the retraction of some synapses based on Hebbian principles and the situation in which the strengths of anatomically fixed synapses are modified based on Hebbian principles?
- The nature of the input to a neural region is an important factor in determining the functional properties of that region. What other factors interact with activity-dependent developmental processes in order to influence the functional properties of a neural region?
- To what extent, and in what ways, is the development of central cortical areas, such as association areas, dependent on the functional properties of peripheral cortical areas?
- From a functional viewpoint, why might a neural system that can be characterized as a collection of relatively independent neural modules be preferable to a system that cannot be characterized in this way?

lack the necessary spatial or temporal resolution and are not suitable for longitudinal studies. In addition, it is relatively easy to distinguish causal relationships from correlational relationships between two variables using computer simulation, whereas it is often difficult to make this distinction in neuroscientific studies. Consequently, computer simulation provides one of the only tools available for evaluating and refining our large-scale theories of the development of functionally specialized neural modules.

Acknowledgements

I thank J. Ringo for commenting on an earlier version of this manuscript. This work was supported by NIH research grant R29-MH54770.

References

- 1 Fodor, J.A. (1983) *The Modularity of Mind*, MIT Press
- 2 Gazzaniga, M.S. (1989) Organization of the human brain *Science* 245, 947-952
- 3 Greenfield, P.M. (1991) Language, tools, and brain: the ontogeny and phylogeny of hierarchically organized sequential behavior *Behav. Brain Sci.* 14, 531-595
- 4 O'Leary, D.D.M. (1989) Do cortical areas emerge from a protocortex? *Trends Neurosci.* 12, 400-406
- 5 Changeux, J.-P. and Danchin, A. (1976) Selective stabilization of the developing synapses *Proc. Natl. Acad. Sci. U. S. A.* 70, 2976-2978
- 6 Elman, J.L. et al. (1996) *Rethinking Innateness*, MIT Press
- 7 Greenough, W.T., Black, J.E. and Wallace, C.S. (1987) Experience and brain development *Child Dev.* 58, 539-559
- 8 Jacobs, R.A. (1997) Nature, nurture, and the development of functional specializations: a computational approach *Psychonomic Bull. Rev.* 4, 299-309
- 9 Karmiloff-Smith, A. (1992) *Beyond Modularity: A Developmental Perspective on Cognitive Science*, MIT Press
- 10 Quartz, S. and Sejnowski, T.J. (1997) The neural basis of cognitive development: a constructivist manifesto *Behav. Brain Sci.* 20, 537-596
- 11 Neville, H.J. (1995) Developmental specificity in neurocognitive development in humans, in *The Cognitive Neurosciences* (Gazzaniga, M.S., ed.), pp. 219-231, MIT Press
- 12 Sur, M., Pallas, S.L. and Roe, A.W. (1990) Cross-modal plasticity in cortical development: differentiation and specification of sensory neocortex *Trends Neurosci.* 13, 227-233
- 13 Chenn, A. et al. (1997) Development of the cerebral cortex: mechanisms controlling cell fate, laminar and areal patterning, and axonal connectivity, in *Molecular and Cellular Approaches to Neural Development* (Cowan, W.M., Jessell, T.M. and Zipursky, S.L., eds), pp. 440-473, Oxford University Press
- 14 Bever, T.G. (1980) Broca and Lashley were right: cerebral dominance is an accident of growth, in *Biology and Language* (Kaplan, D. and Chomsky, N., eds), pp. 186-230, MIT Press
- 15 Geschwind, N. and Galaburda, A.M. (1987) *Cerebral Lateralization: Biological Mechanisms, Associations and Pathology*, MIT Press
- 16 Kosslyn, S.M. (1987) Seeing and imagining in the cerebral hemispheres: a computational approach *Psychol. Rev.* 94, 148-175
- 17 Jacobs, R.A. et al. (1991) Adaptive mixtures of local experts *Neural Comput.* 3, 79-87
- 18 Barlow, H.B. (1986) Why have multiple cortical areas? *Vis. Res.* 26, 81-90
- 19 Cowey, A. (1981) Why are there so many visual areas? in *The Organization of the Cerebral Cortex* (Schmidt, F.O., Warden, F.G., Edelman, G. and Dennis, S.G., eds), pp. 395-413, MIT Press
- 20 Jacobs, R.A. and Jordan, M.I. (1992) Computational consequences of a bias toward short connections *J. Cogn. Neurosci.* 4, 323-336
- 21 Shrager, J. and Johnson, M.H. (1996) Dynamic plasticity influences the emergence of function in a simple cortical array *Neural Netw.* 9, 1119-1129
- 22 Jordan, M.I. and Jacobs, R.A. (1994) Hierarchical mixtures of experts and the EM algorithm *Neural Comput.* 6, 181-214
- 23 Peng, F., Jacobs, R.A. and Tanner, M.A. (1996) Bayesian inference in mixtures-of-experts and hierarchical mixtures-of-experts models with an application to speech recognition *J. Am. Stat. Assoc.* 91, 953-960
- 24 Jacobs, R.A. and Kosslyn, S.M. (1994) Encoding shape and spatial relations: the role of receptive field size in coordinating complementary representations *Cognit. Sci.* 18, 361-386
- 25 Marsolek, C.J. (1995) Abstract visual-form representations in the left cerebral hemisphere *J. Exp. Psychol. Hum. Percept. Perform.* 21, 375-386
- 26 Marsolek, C.J., Schacter, D.L. and Nicholas, C.D. (1996) Form-specific visual priming for new associations in the right cerebral hemisphere *Mem. Cognit.* 24, 539-556
- 27 Laeng, B. (1994) Lateralization of categorical and coordinate spatial functions: a study of unilateral stroke patients *J. Cogn. Neurosci.* 6, 189-203
- 28 Kosslyn, S.M. et al. (1992) Categorical versus coordinate spatial representations: computational analyses and computer simulations *J. Exp. Psychol. Hum. Percept. Perform.* 18, 562-577
- 29 Milner, P.M. (1974) A model for visual shape recognition *Psychol. Rev.* 81, 521-535
- 30 Ebbesson, S.O.E. (1984) Evolution and ontogeny of neural circuits *Behav. Brain Sci.* 7, 321-366
- 31 Johnson, M.H. and Vecera, S.P. (1996) Cortical differentiation and neurocognitive development: the parcellation conjecture *Behav. Process.* 36, 195-212
- 32 Huttenlocher, P.R. (1990) Morphometric study of human cerebral cortex development *Neuropsychologia* 28, 517-527
- 33 Purves, D. (1986) The trophic theory of neural connections *Trends Neurosci.* 9, 486-489
- 34 Cowan, W.M. et al. (1984) Regressive events in neurogenesis *Science* 225, 1258-1265
- 35 Purves, D. and Lichtman, J.W. (1985) *Principles of Neural Development*, Sinauer
- 36 Miller, K.D. (1998) Equivalence of a sprouting-and-retraction model and correlation-based plasticity models of neural development *Neural Comput.* 10, 529-547
- 37 Hinton, G.E. (1989) Connectionist learning procedures *Artif. Intell.* 40, 185-234
- 38 Minsky, M. (1986) *The Society of Mind*, Simon & Schuster
- 39 Harwerth, R.S. et al. (1986) Multiple sensitive periods in the development of the primate visual system *Science* 232, 235-238
- 40 Kerszberg, M., Dehaene, S. and Changeux, J.-P. (1992) Stabilization of complex input-output functions in neural clusters formed by synapse selection *Neural Netw.* 5, 403-413

Coming soon to *Trends in Cognitive Sciences*

- **Modular function in the vegetative state**, Comment by N.D. Schiff, R. and F. Plum
- **Response from A. Owen and D. Menon**
- **Amygdala circuitry in attentional and representational processes**, by P.C. Holland and M. Gallagher
- **Episodic memory: what can animals remember about their past?**, by D. Griffiths, A. Dickinson and N. Clayton
- **Ten years of the rational analysis of cognition**, by N. Chater and M. Oaksford
- **The constructivist brain**, by S.R. Quartz