The computational modelling of sensitive periods

Michael S. C. Thomas and Mark H. Johnson

School of Psychology, Birkbeck College, London, UK

Running head: Modelling of sensitive periods

Address for correspondence:

Dr. Michael Thomas
Developmental Neurocognition Laboratory
School of Psychology
Birkbeck College, University of London
Malet Street, Bloomsbury
London WC1E 7HX, UK
Email: m.thomas@bbk.ac.uk
Tel.: +44 (0)20 7631 6386
Fax: +44 (0)20 7631 6312

Typesetting note: We have used Greek symbols in the text but we have found these sometimes do not appear if the document is opened on a computer loaded with different fonts. The symbols are Δ (upper case delta) and ε (epsilon). These symbols appear on p.3, paragraph 4, line 2 (below equation 1): "where a_i is the activation of the sending unit and a_j is the activation of the receiving unit, w_{ij} is the connection strength between them, [*uppercase delta*] is the change in this strength, and [*epsilon*] is the 'learning rate parameter' (see, e.g., O'Reilly and Munakata, 2000, equation 4.2)."

Introduction

In the continuing debate on sensitive periods, Tyler (this issue) argues for a mechanistic explanation of sensitive periods in development, rather than simply deriving a relationship between plasticity and age. Armstrong et al. (this issue) endorse convergent approaches to assessing types of plasticity, including the use of behavioural evidence, neurophysiological evidence, functional magnetic resonance imaging, event related potentials, and an appeal to evolutionary perspectives. In this contribution, we propose that a computational level of analysis is a key component in understanding the mechanisms through which functional plasticity alters in the cognitive system. To support our case, we discuss three examples of specific computational models that exhibit reductions in plasticity, and show how these models relate to Johnson's (2005) three proposals for the ways in which sensitive periods might end: *endogenous*, *self-terminating*, and *stabilisation*. Typically, we will find that implemented computational models of sensitive periods demonstrate multiple influences at work when functional plasticity reduces. Further, we suggest that computational modelling will allow us to understand how different factors interact to result in a functional reduction of plasticity in different cases.

The importance of computational implementation

Implementation serves to evaluate the assumptions contained within a theoretical proposal. It may be as straightforward as demonstrating that, in a given cognitive domain, turning down a 'learning rate' parameter in a model of development is sufficient to capture the behavioural data indicating a sensitive period. Models are a concrete way to ask, does the theory really work? However, more often multiple assumptions are contained within any theory, and models serve as an exploration and explication of how these factors may interact in driving the functional plasticity of a system. Further, models may generate novel, testable predictions for how plasticity can be increased or decreased in the system. Most importantly, implementation forces the modeller to make decisions about hidden assumptions within verbally specified theories.

Three issues come to the fore when taking a computational perspective regarding sensitive periods in functional brain development: (1) What is the actual nature of the *representations* used to encode the problem domain? It turns out that both the overlap between the representations generated by old and new experiences, and the systematicity within problem domains can both be influential in determining functional plasticity. (2) What is the *frequency* with which the system encounters various experiences? It turns out that under some conditions, frequency can overcome changing conditions of internal plasticity. (3) What level of *processing resources* is available to the system? It turns out that under some conditions, changes in resources can be directly equivalent to changes in plasticity, particularly in parallel processing systems, and further that competition for limited resources can account for many instances of reduced functional plasticity.

Let us consider the last of these three points. Processing resources are of particular relevance where recovery from damage is used as a metric of plasticity. Evidence of 'crowding effects' in children who have suffered brain damage indicates that capacity limitations can influence cognitive development (Anderson et al., 2001). A crowding effect describes the situation where after recovery, there is a generalised depression of neuropsychological functions rather than specific cognitive deficits, as if the remaining system has the computational properties but not the capacity to follow the normal course of development. It has been argued that children's ability to recover from brain damage depends to some extent on their pre-morbid level of processing resources, termed cerebral or cognitive 'reserve' (Dennis, 2000; Stern, 2002). The greater the pre-morbid level of resources, the better the prospect for recovery. A focus on resources prompts the following conclusion: *one cannot interpret a developmental failure to recover from brain damage as a lower level of plasticity unless it is established that the domain(s) in question can definitely be acquired with the reduced level of resources*, were this reduced level to be present at the start of development. Thus, when de Schonen et al. (2005) observe in children with pre-, peri-, or post-natal brain damage a failure to later acquire face recognition expertise, the authors interpret this in terms of 'poor postlesional face-processing plasticity' (p. 184); yet it may be that the remaining processing resources available to the child were simply insufficient to acquire the normal level of expertise whatever the level of plasticity.

Alternative explanations of this nature derive from the requirement to make decisions about resources when building a model. Implementation, for example, would force a modeller to make a decision about what is happening inside a learning system during a period of sensory deprivation. However, the fact that models of development employ analytically derived learning algorithms itself leads to new candidate explanations of changes in functional plasticity. Take the well-known example of Hebbian learning. Within the brain, Hebbian learning can be grossly characterised as ‘cells that fire together, wire together’. More specifically, the change in the connection strength between two neurons is held to be proportional to the product of their correlated activity. More formally,

$$\Delta w_{ij} = \varepsilon a_i a_j \qquad (1)$$

where a_i is the activation of the sending unit and a_j is the activation of the receiving unit, w_{ij} is the connection strength between them, Δ is the change in strength, and ε is the ‘learning rate parameter’ (see, e.g., O’Reilly and Munakata, 2000, equation 4.2). The learning rate parameter is employed when multiple associations are to be learnt in the same network. Its value is typically set at less than 1 to prevent wild oscillations between different connection strengths after each training experience and instead encourage the network to converge on a compromise value that will accommodate all associations. Clearly, the plasticity of a system using this algorithm can be manipulated just by altering the ‘learning rate parameter’. But less obviously, increases in the activation of either the sending or receiving unit themselves increase plasticity. That is, under the terms of the Hebbian algorithm, simply *a more activated system will be more plastic one.*[1] It is not clear whether this candidate mechanism for altering plasticity has relevance for brain development. Event-related potential studies of brain activity indicate that voltage potentials are of greater amplitude earlier in development (see, e.g., Nelson & Monk, 2001, Figure 9.5), though other factors such as skull thickness and conductivity may partially explain this. Brain metabolism measured through PET shows a rising then falling profile across development, with a peak in mid-childhood, though synaptic density appears to peak around 1 year of age (Chugani et al., 1987; Huttenlocher, 2002). In fMRI, the BOLD response in children and adolescents appears to be similar to that in adults in time course and peak amplitude (Casey, Davidson, & Rosen, 2002), although on individual tasks, brain activations in children have been found to be more widespread than in adults (e.g., Casey et al., 1997). The extent to which these neurophysiological measurements relate to the working computational learning algorithm in the brain, and their changes during development, may be a promising novel line of enquiry in developmental cognitive neuroscience.

---

[1] See Mareschal and Bremner (2006) for an application of this idea to infant behavioural development.

We now turn to some examples of implemented models, where the impact of factors such as representational overlap, frequency, and resource level becomes apparent. First, let us recap Johnson's (2005) three classes of explanation for the end of sensitive periods. These are that (a) the termination arises from endogenous factors controlled by maturation or an external environmental "trigger", (b) learning is self-terminating, in that the system drives itself into a representational state where it is no longer responsive, and (c) underlying plasticity does not actually reduce but the constraints on plasticity (such as environmental inputs) become stable. The following three examples all exhibit sensitive periods that come to an end, and each appeals to one of the above explanations. Note that all examples will use algorithms that contain a 'learning rate' parameter but in all cases, that parameter is held constant throughout training.

Example 1: Chick imprinting and the self-terminating sensitive period

O'Reilly and Johnson (1994) constructed a model of filial imprinting in the chick brain. When chicks are exposed to visual stimuli early in life, they can develop a strong preference for a given object. This imprinting can only be established in a specific period of life, is relatively unaffected by subsequent exposure to different objects, and is self-terminating in that the sensitive period is experience driven rather than based on strict chronological age. O'Reilly and Johnson's (1994) neurocomputational model was based on the known neuroanatomy of the chick forebrain and contained several features, including the development of translation invariance for objects presented on its retina. Here we will just concentrate on how its representations developed, simplifying the dynamics of the model somewhat. The model was self-organising, in that it developed representations on an output layer based on exposure to patterns presented on an input layer. In the simulations capturing the closing of the sensitive period, the model was trained on Object A for 100 presentations. It was then trained on an entirely dissimilar Object D. After 150 presentations of D, the network switched its preference from A to D, where preference was assessed by the total activation on the output layer produced by each object. However, if the model was initially trained for only 25 presentations longer (125 presentations of Object A), its preference did not switch to D even after 900 presentations of Object D. Experience-dependent self-organisation led to the closing of the sensitive period at 125 presentations of A. This provides an example of how self-termination of plasticity might work.

It is instructive to consider how this process worked in terms of underlying computations. Increased training on Object A led to further recruitment of units on the output layer to represent this input pattern. After 125 presentations of Object A, the majority of units on this layer were now representing Object A. Since Object D was dissimilar to A (their representations were non-overlapping), it could only activate and therefore attempt to recruit different output units to those activated by A. That is, it could not impinge on the units already recruited by A due to the lack of similarity. As a result, however much learning took place on D, there only remained a minority of the output units that could become selective for this stimulus. Given that the model's stimulus preference was driven by total activation engendered on the output layer, D could never become the preferred stimulus once A had recruited a majority of the output units. There were insufficient resources left to permit this (see O'Reilly & Johnson, 1994, p.374).

Therefore, although this is clearly an instantiation of a *self-terminating* sensitive period, it arises due to *competition for limited resources* and a *lack of representational overlap between new and old experiences* in this implementation.

Example 2: Non-native phoneme discrimination and the sensitive period ended by stabilisation

Monolingual Japanese speakers have difficulty discriminating the English /r/ and /l/ sounds despite repeated exposure to words containing them, consistent with reduced functional plasticity for the acquisition of non-native phonemic contrasts in second language learners. However, if exaggerated versions of /r/ and /l/ phonemes are presented to monolingual Japanese speakers, they can learn to distinguish both these phonemes and subsequently normal exemplars of the /r/ and /l/ phonemes (McCandliss, Fiez, Protopapas, Conway, & McClelland, 2002). McClelland, Thomas, McCandliss and Fiez (1999) constructed a neurocomputational model to explore how this reduction in plasticity might take place in monolingual speakers. The model used a self-organising architecture, with an input layer on which the phonemes were presented and an output layer that had to develop the relevant categories. Two versions of the model were trained. A 'Japanese' model learned a single category of phonemes in the /l/-/r/ region of input space and learned a single output category, while an ‘English’ model was presented with two partially overlapping input categories standing for tokens of /l/ and /r/ and learned two output categories. In the transfer condition, ‘adult Japanese’ networks with 300 epochs of training were exposed to the English-like environment with separate /l/ and /r/ tokens. None subsequently reorganised their output layer into two output categories. However, when ‘exaggerated’ tokens of /l/ and /r/ were used for the two input categories, all ‘adult Japanese’ networks learned to discriminate these stimuli within only a few epochs of their introduction into the training set and this discrimination then extended to the original exemplars.

Again, it is instructive to consider the exact representations used. Each phoneme was represented by a 3x3 square on a grid-like input layer. The single ‘Japanese’ /l/-/r/ input was a 3x3 square in the centre of the input layer. After training, a single output category came to represent this input pattern. The ‘English’ /l/ and /r/ categories were represented by two 3x3 squares on the input layer that overlapped by one row. Their representations had 3 squares in common and 6 squares separate. In the ‘English’ condition, the 6 non-overlapping squares were sufficient to drive the development of two separate output categories. When the ‘Japanese’ net was exposed to ‘English’ input, the two ‘English’ phoneme categories overlapped the single ‘Japanese’ category by two rows each, i.e., each shared 6 squares with the single ‘Japanese’ category and differed by only 3.

Consider a trained ‘Japanese’ network with its one output category. It is now presented with the two novel ‘English’ input categories. When either novel input is presented, the network receives activation from 6 squares that fall within its original input category and only 3 that fall outside. The final output state is the result of a competition, in which the 6 old inputs defeat the 3 new: the novel input is assimilated to the original single category, and the network does not register that it has seen something new. In order for plastic change to occur, new units must win the competition on the output layer. The exaggerated tokens of the ‘English’ /l/ and /r/ categories are created so that they only overlap with the single ‘Japanese’ phoneme by 3 squares; 6 squares fall outside the old category. Now the network receives signals from 6 squares that the input is something new and only 3 that it is old. Different units win the competition to become activate on the output layer, and this causes reorganisation into two output categories. These categories can then also be activated by the original /l/ and /r/ tokens, since these overlap their exaggerated versions by 6 squares.

In this model, then, the sensitive period of the self-organisation ended because its input had stabilised. Although different tokens appeared in its environment corresponding to the shift to ‘English’ input, the representational overlap between old and new experiences was so great that the learning system was essentially “blind” to the change. Only when the difference between exemplars was artificially increased was the latent plasticity of the system revealed and reorganisation triggered. Here is an example of Johnson’s *stabilisation* class of termination, but one that crucially depends on *representational overlap* for its implementation.

Example 3: The emergence of specialised functional structure and the sensitive period ended by endogenous factors

The preceding examples have focused on sensitive period effects in self-organising systems. Research has also explored sensitive periods in associative systems that are required to learn input-output mappings. These have included research on sensitive periods for recovery from damage (Marchman, 1997) and age-of-acquisition effects (Ellis & Lambon Ralph, 2000; Lambon Ralph & Ehsan, in press), both in the domain of language acquisition. In this section, we briefly discuss some results from our own simulation work extending the findings of Marchman (1997).

Marchman (1997) employed the English past tense as a test domain to study acquisition, loss, and recovery in associative networks. The English past tense is of note because it is characterised by a predominant rule (e.g., talk-talked, drop-dropped, etc.) that extends to novel stems (e.g., wug-wugged), but also contains exception verbs (go-went, hit-hit, sing-sang). This aspect of grammar has been much studied because of the problems its dual regular/irregular structure presents for children during language acquisition. It has even been proposed that different brain areas become specialised for the processing of regular and irregular verbs (see, e.g., Tyler, Marslen-Wilson & Stamatakis, 2005). The English past tense is of interest here because it is possible to simulate the emergent specialisation of regular and irregular verbs to different pathways in an associative network (Thomas & Karmiloff-Smith, 2002; Thomas & Richardson, 2006). The problem can therefore additionally serve as a test domain with which to explore sensitive periods in the emergence of specialised functional structure. This issue is important because plenty of evidence suggests that children suffering unilateral brain damage can reorganise their systems to achieve a functional structure sufficient to generate behaviour in the normal range, while adults who suffer similar damage exhibit persisting deficits. Aphasia after left hemisphere damage is one example (see Bates & Roe, 2001). Such evidence implies a sensitive period for when functional structures can be reorganised after damage.

Our simulations used an associative network with two pathways, trained using the backpropagation algorithm. The architecture is shown in Figure 1. The input layer is connected to output layer either directly or via a layer of intermediate processing units. During training, the direct route is more suited to learning regular past tenses and the general rule, while the indirect route comes to specialise in exception mappings that require its additional computational power (see Thomas & Karmiloff-Smith, 2002, for details of this model). We assessed the functional plasticity of this system by measuring its recovery from damage at different points in training. In the normal condition, a network is trained for 500 epochs. A lesion occurring at 490 epochs would only therefore give the network 10 epochs to recover. The confound of lesion age and recovery time can of course suggest poorer plasticity later in training. However, we can control for this artefact by extending training beyond the normal period, so that each network has 500 epochs to recover from damage irrespective of when the damage occurred. The network was damaged either prior to training, or after 10, 50, 100, 250, 400, 450, or 490 epochs of training by lesioning 75% of the connections in both pathways. Its ability to recover was then assessed.

<Insert Fig.1 about here>

Crucially, the network was also given an *endogenous reduction* in its plasticity. From 100 epochs onwards, any network connection below a given threshold had a small probability of being pruned away (i.e., set to zero for the remainder of training), implementing the idea that the network is initially over-resourced but then prunes away unnecessary connections (Huttenlocher, 2002). With fewer connections, the network's ability to learn is reduced. The 100-epoch onset presumed an endogenous trigger for pruning in the model.

Figure 2 shows the normal endstate performance (grey bars) for the regulars, rule generalisation, and three types of exception verb (labelled EP1, EP2, EP3f). It also demonstrates the level of endstate recovery achieved following damage at different points during training. Performance levels are shown both for recovery at the completion of 500 epochs, where later lesions will have had shorter recovery times (white bars), and following a fixed recovery period of 500 epochs post lesion (black bars). Regulars and rules indicated little evidence of sensitive periods in this associative system, with similar levels of recovery whenever the damage occurred. Regular patterns and rule generalisation retained their functional plasticity because of the high type frequency and systematicity amongst regular past tenses in the training set (see Lambon Ralph & Ehsan, in press, and Seidenberg & Zevin, 2006, for discussions of the influence of systematicity and frequency on age-of-acquisition effects). Regulars are best positioned to use the remaining resources after damage. By contrast, all three types of exception pattern exhibited sensitive periods. In the case of EP3f exception patterns, the sensitive period declined in a roughly linear fashion. These verbs have arbitrary input-output mappings but high token frequency in the training set, and their high token frequency allows the best recovery of the exception patterns. For EP1 and EP2 exception patterns, the decline in recovery with age was steeper; perhaps one might call these 'critical' rather than sensitive periods. Overall, the results show that *within the same architecture, sensitive and critical periods can appear in some parts of the problem domain but not others*, depending on the nature of the mapping problem and on frequency effects.

<Insert Fig.2 about here>

Figure 3 plots the proportion of connections remaining in one of the pathways of the network and depicts the gradual reduction through pruning as well as the sudden drop after lesion is applied at an early and a late point in training. Importantly, although pruning was an endogenous process, it was also influenced by activity-dependent changes in the network. When a lesion occurred early in training, the network was able to take advantage of the remaining resources and *fewer connections were pruned*.

It turned out that both resources and pruning were key in generating the sensitive periods observed in this model. When the model was trained with fewer resources (units) in the indirect pathway, sensitive periods appeared for all pattern classes. When the normal network was trained without pruning, none of the pattern classes exhibited sensitive periods.

<Insert Fig.3 about here>

Figure 4(a) shows the relative functional specialisation of each pattern type to the direct (+ve) or indirect (-ve) pathways of the associative network. Figure 4(b) focuses on one specific contrast in emergent specialisation, rule formation versus EP3f patterns, assessed across a fixed period of recovery after damage. In the normal condition, regular and rule generalisation revealed partial specialisation to the direct pathway, while exception patterns showed differing degrees of specialisation to the indirect pathway. If both the routes of the network were damaged prior to training, this immediately changed how each pattern class used the two pathways. The indirect pathway was relied upon more heavily. However, as damage occurred later in training, this pattern progressively changed, with increasing reliance on the direct pathway to drive recovery (even when recovery time was controlled). Two points are of note: first, we see here sensitive periods for the emergence of specialised functional structure, with different functional structures arising depending on the time of damage. Second, for the exception verbs, the sensitive period for functional structure corresponded with a sensitive period for behaviour (i.e., the alternate functional structure was less able to support recovery); but for regular verbs and rule generalisation, the sensitive period for functional structure had no corresponding sensitive period in behaviour. *The sensitive periods for functional structure and for behaviour could therefore dissociate*.

<Insert Fig.4 about here>

The explanation for these effects involves several factors. Broadly, the results depend firstly on how well different pattern types can exploit the resources remaining at different points in training, based on their frequency and similarity. Later lesions cause more reduction in resources because they come on top of losses through pruning. Early damage can retard the endogenous pruning process. Systematicity, high type frequency, and high token frequency, all advantage a pattern class in making use of remaining resources. Secondly, the two pathways have different plasticity at an algorithmic level. It takes more training to alter the two sets of connections arranged in series in the indirect route than it takes to alter the single set in the direct route. Thirdly, later in training, connections in each pathway become larger and if these connections are not useful for driving behaviour after damage, they take longer to reset (an effect called 'entrenchment'). These three factors interact to determine which pattern classes will recover and how the two pathways will be used.

This simulation is useful because it can begin to explore the relationship (and possible mismatch) between sensitive periods in behaviour and in the emergence of specialised functional structure, but once more, the effects were mediated by similarity, frequency, and resources. In terms of Johnson's (2005) proposals for how sensitive periods end, this model implemented an *endogenous* process of pruning. Yet even these endogenous factors *interacted with activity dependent processes* in fashioning the final shape of the sensitive periods in plasticity.

Conclusion

We began by endorsing the importance of specifying the underlying computational mechanisms of plasticity change in order to turn descriptions of sensitive periods into explanations, and by arguing for the utility of implemented neurocomputational models in this endeavour. Implementation forces clarity, reveals hidden assumptions, and generates new candidate explanations and testable hypotheses. In three examples, we illustrated implementations of Johnson's (2005) proposals for how sensitive periods might end. In each case, implementation demonstrated multiple additional factors at play that interacted with the closing of sensitive periods, including the similarity between representations, the frequency with which certain experiences occurred, and resource levels within the system. We believe that discovery of the full repertoire of mechanisms through which functional plasticity is modulated must rely on a programme of computational modelling integrated within the multidisciplinary exploration of sensitive periods in development.

Acknowledgements

This research was supported by MRC Career Establishment Grant G0300188 to Michael Thomas.

References

Anderson, V., Northam, E., Hendy, J., & Wrennall, J. (2001). *Developmental neuropsychology: A clinical approach*. Psychology Press: East Sussex, UK.

Armstrong, V. L., Brunet, P. M., He, C., Nishimura, M., Poole, H. L., & Spector, F. J. (xxxx). What is so Critical? A commentary on the Re-Examination of Critical Periods. *Developmental Psychobiology*

Bates, E., & Roe, K. (2001). Language development in children with unilateral brain injury. In C. A. Nelson & M. Luciana (Eds.), *Handbook of developmental cognitive neuroscience* (p. 281-307). Cambridge, Mass: MIT Press.

Casey, B. J., Davidson, M., & Rosen, B. (2002). Functional magnetic resonance imaging: basic principles of application to developmental science. Developmental Science, 5:3, 301-309.

Casey, B. J., Trainor, R. J., Orendi, J. L., Schubert, A. B., Nystrom, L. E., Giedd, J. N., Castellanos, F. X., Haxby, J. V., Noll, D. C., Cohen, J. D., Forman, S. D., Dahl, R. E., & Rapoport, J. L. (1997). A developmental functional MRI study of prefrontal activation during performance of a go-no-go task. *Journal of Cognitive Neuroscience*, *9(6)*, 835-847.

Chugani, H. T., Phelps, M. E. & Mazziotta, J. C. (1987). Positron emission tomography study of human brain functional development. *Annals of Neurology*, *22*, 487-497.

Ellis, A. W., & Lambon Ralph, M. A. (2000). Age of acquisition effects in adult lexical processing reflect loss of plasticity in maturing systems: Insights from connectionist networks. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *2*6, 1103-1123.

de Schonen, S., Mancini, J., Camps, R., Maes, E., & Laurent, A. (2005). Early brain lesions and face-processing development. *Developmental Psychobiology*, *46(3)*, 184-208.

Dennis, M. (2000). Childhood medical disorders and cognitive impairment: Biological risk, time, development, and reserve. In: Yeates K. O., Ris, M. D., Taylor, H. G. (Eds.), *Pediatric neuropsychology: Research theory and practice*, (pp. 3-22). New York: Guilford.

Huttenlocher, P.R. (2002). Morphometric study of human cerebral cortex development. In: Johnson, M.H., Munakata, Y., Gilmore, R. (Eds.), Brain Development and Cognition: A Reader (2nd Edition) (pp. 117-128) Oxford: Blackwell.

Johnson, M. H. (2005). Sensitive periods in functional brain development: Problems and prospects. *Developmental Psychobiology*, *46(3)*, 287-292.

Lambon Ralph, M. A., & Ehsan, S. (in press). Age of acquisition effects depend on the mapping between representations and the frequency of occurrence: Empirical and computational evidence. *Visual Cognition*.

Marchman, V. A. (1997). Constraints on plasticity in a connectionist model of English past tense. *Journal of Cognitive Neuroscience*, *5:2*, 215-234.

Mareschal, D., & Bremner, A. J. (2006). When do 4-month-olds remember the "what" and "where" of hidden objects? In: Y. Munakata & M. Johnson (Eds.), *Attention and Performance XXI: Processes of Change in Brain and Cognitive Development*. Oxford: Oxford University Press.

McCandliss, B. D, Fiez, J. A., Protopapas, A., Conway, M., & McClelland, J. L. (2002). Success and failure in teaching the [r]-[l] contrast to Japanese adults: Tests of a Hebbian model of plasticity and stabilization in spoken language perception. *Cognitive, Affective, and Behavioral Neuroscience*, *2(2)*, 89-108.

McClelland, J. L., Thomas, A. G., McCandliss, B. D., & Fiez, J. A. (1999). Understanding failures of learning: Hebbian learning, competition for representational space, and some preliminary experimental data. In J. A. Reggia, E. Ruppin, & D. Glanzman (Eds.), *Disorders of brain, behavior, and cognition: The neurocomputational perspective*. Elsevier: Oxford. Pp75-80

Nelson, C. A., & Monk, C. S. (2001). The use of event-related potentials in the study of cognitive development. In C. A. Nelson and M. Luciana (Eds.), *The Handbook of developmental cognitive neuroscience* (p.125-147). MIT Press: Cambridge, MA.

O'Reilly, R. C., & Johnson, M. H. (1994). Object recognition and sensitive periods: A computational analysis of visual imprinting. *Neural Computation*, *6*, 357-389.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT Press: Cambridge, Mass.

Seidenberg, M. S., & Zevin, J. D. (2006). Connectionist models in developmental cognitive neuroscience: Critical Periods and the paradox of success. In: Y. Munakata & M. Johnson (Eds.), *Attention & Performance XXI: Processes of Change in Brain and Cognitive Development*. Oxford: Oxford University Press

Stern, Y. (2002). What is cognitive reserve? Theory and research application of the reserve concept. *Journal of the International Neuropsychological Society*, *8*, 448-460.

Thomas, M. S. C. & Karmiloff-Smith, A. (2002). Are developmental disorders like cases of adult brain damage? Implications from connectionist modelling. *Behavioral and Brain Sciences*, *Vol. 25*, *No. 6*, 727-788.

Thomas, M. S. C., & Richardson, F. (2006). Atypical representational change: Conditions for the emergence of atypical modularity. In: Y. Munakata & M. Johnson (Eds.), *Attention and Performance XXI: Processes of Change in Brain and Cognitive Development*. Oxford: Oxford University Press.

Tyler, L. K., Marslen-Wilson, W. D., & Stamatakis, E. A. (2005). Differentiating lexical form, meaning, and structure in the neural language system. *Proceedings of the National Academy of Sciences*, *Vol. 102*, *No. 23*, 8375-8380.

Tyler, A. N. (xxxx). When is a description not an explanation? *Developmental Psychobiology*

Figure captions

Figure 1. Architecture of the associative network trained on the English past tense problem. Rectangles represent layers of simple neuron-like processing units, and black arrows represent matrices of connections between layers. Verbs stems were coded on the input layer and past tenses on the output layer using phonological features.

Figure 2. Performance of the network at the end of training (500 epochs) for five pattern classes within past tense: *Regular* (e.g., talk-talked), *Rule* (wug-wugged), *EP1* (hit-hit), *EP2* (sing-sang), and *EP3f* (go-went). The *EP* numbers mark increasing degrees of inconsistency with regular mappings and the *f* registers the high token frequency of this class. Grey bars show normal performance. Black bars show recovery after lesions at different points in training (0, 10, 50, 100, 250, 400, 450, and 490 epochs) with a fixed period of 500 epochs of training post-lesion. White bars show the recovered level of performance at the end of normal training (e.g., a lesion at 490 epochs will have only 10 epochs of training post lesion). Error bars depict standard errors over 6 replications with different initial random seeds.

Figure 3. The proportion of connections remaining in the direct pathway with the combined effects of pruning (onset 100 epochs) and lesions, shown for the normal case, and for lesions after 10 lesion or 400 epochs of training. Similar functions were found for connections in the indirect pathway. (A connection was pruned with 5% probability each epoch if its absolute value was less than 0.5. Lesions probabilistically removed 75% of connections in both pathways.)

Figure 4. (a) Relative specialisation of each pattern class to the direct (+ve) or indirect (-ve) pathway at the end of training, for the normal network and networks recovering from damage at different points in training. Specialisation was assessed using the dissociation methodology of traditional cognitive neuropsychology. [If a pattern class is more specialised to the direct than indirect pathway, it should show a bigger deficit when the direct pathway experiences a further lesion than when the indirect pathway is similarly lesioned. Figure 4 shows the difference in the size of the deficit for each pathway (see Thomas & Karmiloff-Smith, 2002, for details)]; (b) A single

comparison drawn from the above data, depicting the relative specialisation of *rule* versus *EP3f* patterns after a fixed recovery period following damage.

Figure 1.

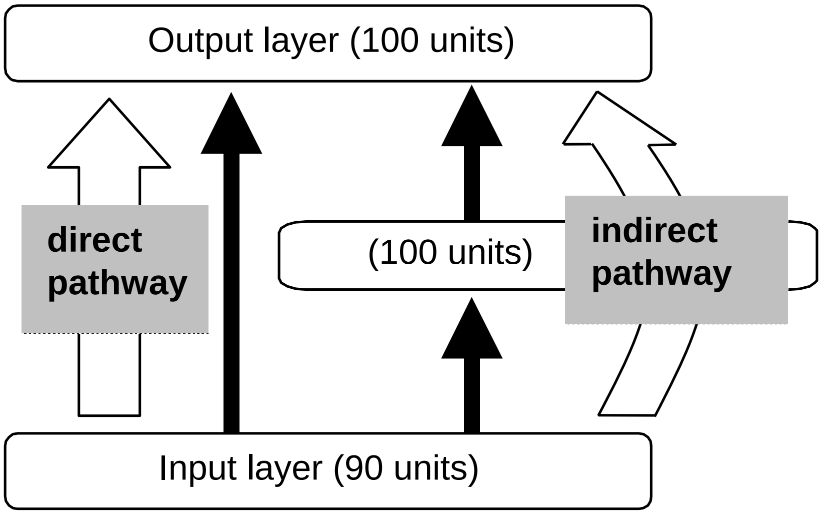
Output layer (100 units)
direct pathway
(100 units)
indirect pathway
Input layer (90 units)

Figure 2.

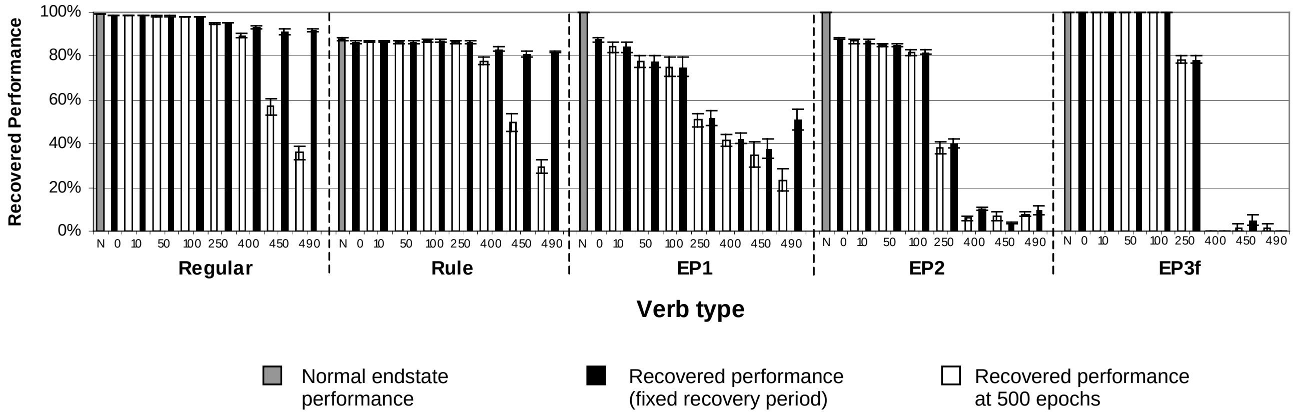
100%
80%
60%
40%
20%
0%
Recovered Performance
N 0 10 50 100 250 400 450 490
Regular
Rule
EP1
EP2
EP3f
Verb type
Normal endstate performance
Recovered performance (fixed recovery period)
Recovered performance at 500 epochs

Figure 3.

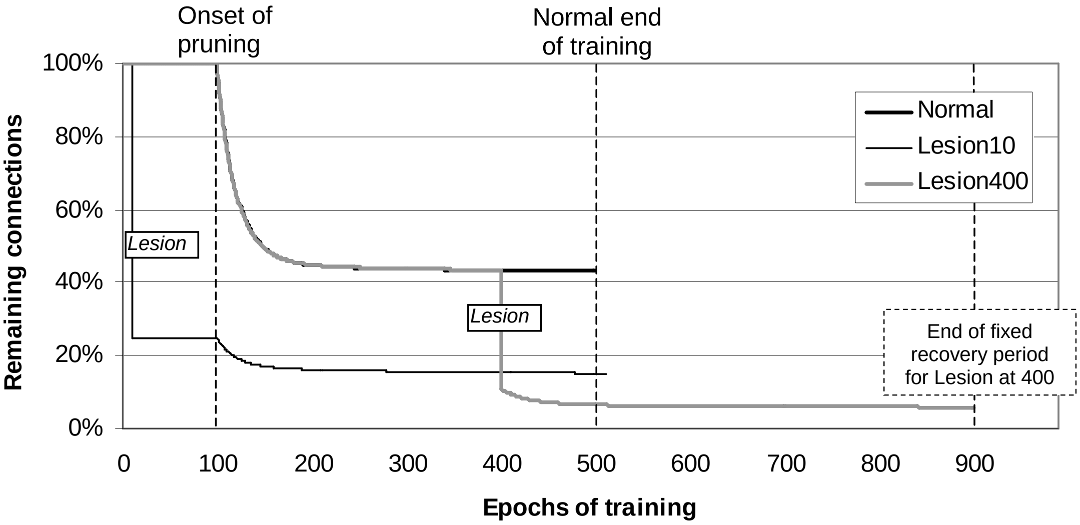
Onset of pruning
Normal end of training
100%
80%
60%
40%
20%
0%
Remaining connections
Lesion
Lesion
Normal
Lesion10
Lesion400
End of fixed recovery period for Lesion at 400
0
100
200
300
400
500
600
700
800
900
Epochs of training

Figure 4a

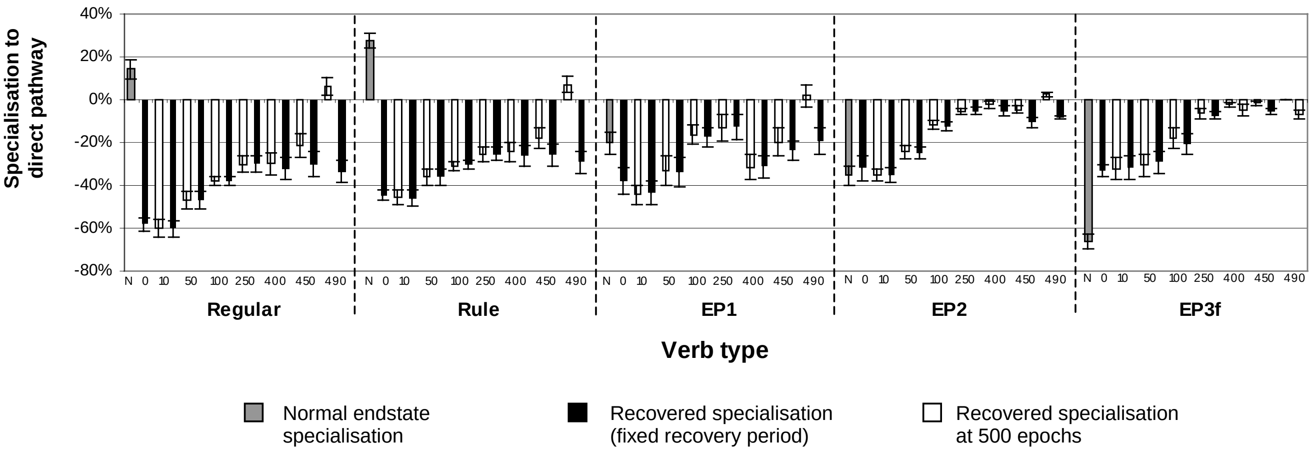

Figure 4b