Cognitive Neuropsychology and Computational Modelling: The Contribution of Computational Neuroscience to Understanding the Mind

Glyn W. Humphreys

Dietmar G. Heinke

and Eun Young YOON

Behavioural Brain Sciences,

School of Psychology,

University of Birmingham,

Birmingham

B15 2TT

Computational neuroscience attempts to use explicit, biologically-inspired, computational models to simulate, predict and explain human performance. In this Chapter, we will review research using this approach to model neuropsychological disorders, particularly disorders of visual object recognition and attention. We argue that computational studies are able to go beyond more traditional 'box and arrows' models of cognition, particularly when disorders result from interactions between different components within a given cognitive system. Such studies provide an important means of helping understand how mental function arises from complex neural networks.

Cognitive neuropsychology: Assumptions and boxology

Cognitive neuropsychological research is concerned with understanding the cognitive disorders suffered by patients after brain lesions, and with using data from such disorders to inform us about the processes that contribute to performance in normal (intact) participants (Coltheart, 1984). Since neuropsychological studies can generate striking, and sometimes counter-intuitive, dissociations between processes, they provide an informative part of current cognitive neuroscience research, adding to those gained from behavioural and imaging procedures carried out with normal individuals (see other Chapters in this volume). For example, it remains the case that some of the strongest evidence that constrains our understanding of the processes involved in normal reading comes from the study of acquired dyslexia. Here there is a double dissociation between phonological dyslexia (impaired reading of nonwords, spared reading of irregular words) and surface dyslexia (which can involve spared reading of nonwords along with impaired reading of irregular words; see Beauvois & Derouesné, 1979; Funnell, 1983; Shallice, Warrington & McCarthy, 1983). This double dissociation has proved difficult to account for using 'single route' models, especially when there are phonological dyslexic patients who seem to show poor comprehension for the irregular words they are able to name, arguing against spared word reading based on access to semantic knowledge (cf. Funnell, 1983; cf. Plaut et al., 1996).However, the neuropsychological data are consistent with a 'dual route' account which distinguishes between a lexical, knowledge-based reading process and a non-lexical, rule-based process, with each syndrome being linked to damage to one of the two routes (phonological dyslexia to damage to the non-lexical

route, surface dyslexia to damage to the lexical process; cf. Coltheart, 1987; Coltheart et al., 1993).

These examples of dyslexia also illustrate two of the fundamental assumptions that are typically made in cognitive neuropsychology, that have aided interpretation of any results (see Caramazza, 1986). One assumption is *modularity.* This assumption holds that a task such as reading can be broken down into a set of component processes (e.g., access to the visual lexicon, access to semantic knowledge, the application on non-lexical spelling-sound rules etc.), and that damage to any of these component processes will lead to an individual having a reduced version of the normal system. Thus, in phonological dyslexia we witness the operation of a lexical reading route without a contribution to pronunciation from spelling-sound rules, while surface dyslexia performance is based on output from the spelling-sound rules without a contribution from the lexical reading route. Inferences can then be made about the nature of the impaired component processes, revealed in the difference between the patient's performance and that of normal participants. For example, we may assume that visual lexical access is frequency-sensitive, given that surface dyslexics may be more likely to read high than low frequency irregular words, in cases where lexical representations for high frequency words are spared (Bub, Cancelliere & Kertesz, 1985). A second assumption typically made is that of *transparency*, that is, that the symptoms that characterize a given neuropsychological disorder are telling either about the underlying component process that has been damaged, or about the spared normal process being used. For instance, we use this assumption when we inferthe regularization errors made by surface dyslexics

when naming irregular words are indicative of the non-lexical rules that determine reading when the lexical route is damaged.

It can also be argued that the idea of modular cognitive systems was greatly encouraged by the adoption of 'box and arrow' theories developed in the 1970's and 80's. Using these models, a given ability would be conceptualized in terms of the representations constructed or accessed at particular stages of processing, along with the processing 'routes' that transform input from one stage to the next (see Figure 1 for an example). Neuropsychological syndromes can be mapped onto such models by positing that a given representation or processing route has been damaged in a particular patient. What was especially noteworthy in the development of these models in 70's and 80's was the attempt to provide a detailed account of the nature of the representations and processing routes involved. This distinguished the emerging work on cognitive neuropsychology from the conceptually similar, but much more simplified, frameworks put forward by the so-called 'diagram makers' at the very start of modern-day neuropsychology (e.g., Lichtheim, 1885). To this day, cognitive neuropsychological accounts of performance, based on these articulated, 'box and arrow' frameworks, remain a powerful influence on neuroscientific thinking.

Figure 1 about here

However, there are important limitations to such theories, and to their ability to account for various neuropsychological disorders where performance appears to emerge from interactions between separable components in the processing system, and for disorders where a given behaviour may be caused by adaptation or a strategy adopted post-lesion. In the next section we discuss two impairments where these problems arise,

that have been the focus of study in our laboratory: optic aphasia (Freund, 1889; Lhermitte & Beauvois, 1973) and visual apraxia (De Renzi, Faglioni & Sorgato, 1982).

Optic aphasia and visual apraxia.

The term optic aphasia is used to describe patients who present with a selective, modality-specific impairment in naming visually presented objects. Naming in other modalities (e.g., to touch, or to a definition) may be relatively intact. In contrast to their poor visual naming, optic aphasic patients can very often perform a gesture to show how an object would be used. This good ability to gesture has frequently been interpreted as evidence for the patient being able to access 'semantic' knowledge about objects, so that the problem is one of name retrieval. However, in at least three studies that have looked in detail at the semantic knowledge such patients do have access to, performance has been found to be impaired. For instance, Riddoch and Humphreys (1987) presented their patient, JB, with sets of 3 objects and asked him to choose which two would be used together (e.g., hammer, nail, spanner). JB performed poorly at this task when the objects were visually presented. This was not due to a lack of semantic knowledge per se, because JB was at ceiling when he was given the names of the objects and asked to make the same choice (see also Caramazza & Hillis, 1995; Yoon, Humphreys & Riddoch, in press). Such results suggest that there was a modality-specific problem in accessing semantics. Despite this, JB was good at gesturing to visually presented objects. How could this be, given that access to semantic knowledge was impaired? There are at least two ways that we can think of this, and one way violates the idea that symptoms are transparent. One account, that does fit the 'classical' modular framework for cognition

(object naming in this instance), assumes that there exist separate 'routes' from visually presented objects to action. A 'semantic' route would involve access to semantic knowledge from vision – much as one would need to access semantic knowledge to access action from an object's name. Alongside a second 'direct' route is assumed, based on associations between the visual representation of an object and an action (see Figure 2; Humphreys & Riddoch, 2003). In optic aphasia, the direct route may continue to operate even if there is a modality-specific problem in accessing semantic knowledge, with the result that gesturing is relatively spared. Alternatively, though, the gestures made by optic aphasic patients may reflect an adaptation on the part of the patient, using information (and strategies) that are not normally involved when we act upon objects. Caramazza and Hills (1995), for example, proposed that such gestures may be derived from partial semantic knowledge that the patient remains able still to access, coupled to a visual problem-solving strategy based on explicitly working out how an object might be used, given its visual properties (see also Hodges et al., 2000). If the latter holds, then the gestures made by such patients are not 'transparent', and do not inform us about the processes normally involved in retrieving actions to objects.

Figure 2 about here

Take now the disorder of visual apraxia, which is difficult to explain in terms of standard modular theories. The label visual apraxia is applied to patients who show impaired actions to visually presented objects, despite having intact recognition and, in some cases, spared object naming (De Renzi et al., 1982; Pilgrim & Humphreys, 1991; Riddoch, Humphreys & Price, 1989). The problem is again modality specific, since the same patient may make a correct gesture when given the names of objects. Now it is

possible to conceptualise this disorder in terms of a modular account of how we recognize and action to objects, based on a 'dual route' account of action (Figure 2). For example, given that object recognition and naming can be spared, we would assume that the semantic would be operating, and that the problem lies within the direct route to action from vision. However, if the semantic route is operating, then it is not clear why this route could not serve to enable the patient to make gestures when objects are presented visually – much as it is assumed that the continued existence of the lexical route for reading enables phonological dyslexic individuals still to read words. Simple reduction of a proposed normal system, here, fails to account for the data.

The Naming and Action Model (NAM).

These problems in the standard assumptions underlying much of cognitive neuropsychology can be addressed, however, once we begin to develop working computational models of performance – particularly, we believe, where the models incorporate some aspects of biological neural systems. Our own work on this topic has focused on the use of artificial neural networks that capture some of the dynamic properties of real neural systems – where whole system performance emerges from local interactions between large numbers of processing units (see McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986). Due to this interactivity, there can be emergent effects on performance that would not be predicted through a standard, modular 'box and arrow' account, even if there is some degree of functional specialisation within a given system. Furthermore, within a dynamic model, the effects of a brain lesion may be not simply to reduce the normal system but change the nature of the interactions between

component processes. This turns out to be important to account for some neuropsychological disorders, such as visual apraxia.

The Naming and Action Model, or NAM, was developed to provide a framework for understanding both normal and disordered abilities both to name and to retrieve actions associated with objects (Yoon, Heinke & Humphreys, 2002). NAM employed a quasi-modular structure, composed of input (perceptual), semantic and output (name and action response) units, arranged in an architecture that conforms to a 'dual route' model of name and action retrieval from objects (see above). This architecture is illustrated in Figure 3. The input to NAM, provided by objects, conformed to a description sensitive to the presence of object features (the number of straight and curved lines, the length and width of the object) in relation to the viewpoint. At a semantic level, units corresponded to the specific item and to its super-ordinate category. At an output level, different sets of units served to categorize the object either in terms of its name (hammer vs. pen) or the action typically performed (e.g., hitting vs. writing). Input was also provided by words, through an orthographic perceptual recognition system. For both words and objects as input, both names and actions could be accessed via the semantic system (the semantically mediated route to naming and to action). In addition, different 'direct' routes were set-up, from words to names (cf. Funnell, 1983) and from structural descriptions for objects through to action. At each level, a winner-take-all selection process operated, with the network 'relaxing' over time, to achieve activation for the output units corresponding to a particular name or category of action. This winner-take-all selection process was based on activation values being passed continuously between local processing units. Hence, though there was structural modularity outputs from the models

were contingent on interactions between modules. This implications for the performance of the model following damage, as we elaborate below.

Figure 3 about here

Consider how actions are selected in NAM. Activation of the action classification units is determined by joint activity coming from both the semantic and direct routes to action. Rapid activation, coming through the direct connections from the structural description system to the action system, 'push' activity states in the action units towards a 'basin of attraction' that would comprise a steady state of activity at this level. This activation then converged with activity arising out of the semantic units, to generate selection of the appropriate action to an object. Now, consider the effects of damaging the connections from the structural descriptions to the action units, so that relatively 'noisy' activity is transmitted along the 'direct' route. Instead of pushing output activity towards a basin of attraction, this noise can move activity away, disrupting activation at an output stage. Importantly, this disruption could occur even if the semantic route to action were operating normally, and even if this semantic activity was sufficient to generate the correct output response in the absence of any 'direct' input – as would be the case if participants were required to make actions to words in the absence of any object. Data on action classification following simulated lesions of both the direct and the semantic routes to action are presented in Figure 4.

Figure 4 about here

Figure 4 here shows that symptoms of either optic aphasia or visual apraxia emerged in NAM, according whether a lesion was applied to the direct route for action (simulating visual apraxia, lesion V→ A, Figure 4), or from visual input into the semantic

system (simulating optic aphasia, lesion V→ S). When there is impaired access to semantics (lesion V→ S), there is impaired visual object naming relative to when there is damage to a direct visual route to action (lesion V→ A), whereas the opposite pattern occurs when the task requires action retrieval (performance is worse after lesion V→ A than lesions V→ S). Visual apraxia emerged from damage to the direct route here because noisy outputs from this route effectively disrupted the activity coming from the semantic routes. Due to the interactivity in the system, there was not mere loss of the direct route, but also a change in the way that semantic activity influenced action selection. However, without any activity generated in the damaged visual route (e.g., when the task was simply to retrieve an action given the name of an object), activation from the semantic route was sufficient to allow action selection to be successful. That is, there was action retrieval from a name but not from an object. This is precisely the pattern of performance that violated the assumption of modularity typically made in cognitive neuropsychology. NAM also simulated optic aphasia, following damage to input coming into the semantic system from vision (lesion V→ S). This simulated lesion led to problems in semantic access from vision, a pattern consistent with the neuropsychological data (Caramazza & Hillis, 1995; Riddoch & Humphreys, 1987). Despite this, actions could be accessed more accurately than names from objects due to the continued operation of the direct route to action. In this case, the direct route provided an early 'push' towards correct categorization in the action system , so that noise within the sematic route to action was not too disruptive for performance.

Now a model such as NAM can provide only an existence proof that a system with this architecture and processing dynamics can generate the two contrasting patterns

of performance characteristic of optic aphasia and visual apraxia. Arguments remain about whether observations such as the good gesturing in optic aphasia reflect the operation of a preserved direct route to action (as suggested by NAM), or whether it reflects a strategic adaptation on the part of patients, where they use visual problem solving combined with partial semantic knowledge to act to objects (i.e., is there a 'non-transparent' pattern of performance?). This kind of argument might be answered by looking in more detail at patient performance – for example, measuring the time taken by a patient to retrieve and enact a gesture, compared with normal participants. Here one might expect any strategic, problem-solving porocess to be relatively slow. However, such predictions are not straight-forward. Thus, although NAM was able to select an appropriate action to an object after lesioning visual access to semantics, the latency of action retrieval was slowed (see Yoon et al., 2002). This NAM too predicts that the latency of action retrieval would be disrupted in optic aphasics. A further way needs to be found to distinguish between these accounts. One way is to look for converging evidence coming from normal participants. For example, it seems unlikely that normal participants would have to use explicit problem-solving strategies to make gestures to objects, given that their ability to act on the basis of semantic knowledge is intact. This also seems especially unlikely when normal participants are required to act rapidly. Rumiati and Humphreys (1998) examined this by having normal participants make gestures to objects or to words, when a fast response deadline was imposed. Relative to when actions were made to words, participants tended to make 'visual' gestures to objects under the deadline conditions (e.g., making an action that was related to the visual properties of the object, even if this did not reflect the object's identity – such as making a writing gesture to a

knife). In contrast, there was an increase in semantic errors when actions were made to words. The tendency to make visual errors when gesturing to objects suggests that participants were acting on the basis of activity rapidly derived from the visual route, while semantic errors arose to words because, for these stimuli, gestures were retrieved on the basis of semantic knowledge. This pattern of visual and semantic gesture errors under response deadline conditions, respectively for objetcs and words, can be simulated by NAM when the activation threshold for responding was decreased (to mimic a fast response deadline)(see Figure 5). These results are consistent with NAM rather than with the suggestion that performance is non-transparent, reflecting strategic adaptation on the part of optic aphasic patients. Perhaps even more importantly, the point is that having an explicit model of performance enables normal as well as abnormal performance to be simulated, so that convergent predictions can be assessed. Notions about transparency can be tested.

Figure 5 about here

Frameworks for multiple dissociations: the Selective Attention and Identification Model (SAIM)

One complaint sometimes made by non-neuropsychologists about using neuropsychological data for building and testing theories is that, in some areas, there seem to be a bewildering number of dissociations – often between single patients – making it appear as if any arbitrary pattern of deficit could arise following a brain lesion. In such cases it is tempting to think that different deficits stem from idiosyncratic learning experiences or the prior interests of the patient, and so are not informative about

the general cognitive architecture of the mind. A particular example here comes from the disorder of unilateral visual neglect, where patients fail to respond to stimuli presented on the contralesionsal side of space following their brain injury (e.g., Heilman, Watson & Valenstein, 1985). Over the past twenty years, there have been numerous case reports of dissociations between the symptoms found in neglect patients (see Halligan et al., 2003, for a review). These dissociations can vary from deficits in near but not far space, or vice versa (e.g., Pitzalis et al., 2001; Pizzamiglio et al., 1989), through to deficits that seem related to the presence of multiple independent objects in the field compared with cases where the positions of parts with respect of a single object seem important (e.g., Humphreys & Heinke, 1998). This last dissociation, apparently between a spatial representation of independent objects and a representation of parts within an object, has even been reported within a single patient, as we elaboarte below.

Humphreys and Riddoch (1994, 1995) examined a patient, JR, who had sustained bilateral damage to his parietal cortex as well as his right cerebellum, following multiple strokes. In an initial test, JR was asked to read words and nonwords positioned randomly on a page. JR's pattern of performance is illustrated in Figure 6. What was interesting was that, when JR mis-identified a particular stimulus, the errors were more pronounced at the left ends of the strings. Also, these mis-identification errors were more pronounced on nonwords than on words (i.e., they appeared to be affected by top-down knowledge). In contrast, he made complete omissions of some stimuli that fell on the right side of the page. These omissions were not affected by the lexical status of the letter string. These apparently opposite forms of neglect did not only occur in reading tasks, but they were also apparent when JR was asked to name pictures on a page. Here he continued to make

right-side ommisions whilst mis-identifications of pictures tended to reflect the right-most features in the objects, with the left-most features being neglected. Thus when identifying single items, JR seemed not to 'weight' the left-side features very strongly, whereas he appeared not to 'weight' the right side of space strongly when scanning around multiple objects on a page. These two deficits may stem, respectively, from the right and left-parietal lesions in JR's case. To make sure that omissions were not made simply because the stimuli fell into a blind area of field, we conducted a further study where, using the same stimuli, we tried to bias JR to code visual elements as a single object or as multiple separate objects. We presented him with words and nonwords in large print, so that they spanned the width of an A4 page. We then had him either try and read the whole stimulus, or to read out the letters making up the stimulus (e.g., reading a word as a word or as 'w', 'o', 'r' and 'd'). When asked to read the whole string, JR made left-side errors, typically mis-identifying the letters present (e.g., LIGHT → 'might'). In contrast, when required to identify each letter in turn, he made right-side ommisions (e.g., LIGHT → 'l', 'i', 'g', 'h'). In this instance, when reading individual letters he omitted letters on the right that he had read formerly, when processing the whole string, whilst be correctly identified letters on the left that previously he had earlier mis-identified. These results indicate that the spatial positions of the elements in the world are less crucial than the way that the elements are represented for the task at hand. There appear to be separate spatial representations of parts within objects, and representations of the relations between independent objects; these different spatial representations are separately affected in patient JR (i.e., there is left neglect of the psace within-objects, and

right neglect of a between-object spatial code). Humphreys and Riddoch (1994, 1995) termed these within- and between-object codes.

Now, how are we to understand the relations between these codes? Our view on this is that it is extremely helpful to try and capture such patterns of results within an explicit model of performance, which then helps conceptualise how different disorders can emerge, perhaps after lesions to different loci in a processing system. The Selective Attention for Identification Model (SAIM) provides an example of this, as it generates an explicit account of how within- and between-object neglect can emerge in a framework in which the two codes are used for specific computational purposes (see Heinke & Humphreys, 2003).

Like NAM, SAIM is a quasi-modular model, involving several modules in structural terms, but with interactions within and across modules generating dynamic and interactive performance. SAIM's architecture is depicted in Figure 6. The aim of the model is to achieve translation-invariant object recognition. It does this by mapping input from any lateral position on the retina into a 'Focus of attention' that has a fixed size based on one object. Recognition units, at the higher-end of SAIM ('template units' in the 'Knowledge network') then respond to the presence of active pixels at particular locations in the FOA, but since these pixels are activated from across the retina, the recognition process is translation invariant. Activation in the FOA is itself controlled by two networks: the 'Contents' net and the 'Selection' net. The Contents net can be thought of as a connection matrix specifying all possible relationships between locations on the retina and locations in the FOA. Here a high level of activation in one unit in the contents network instantiates a particular correspondence between a given retinal location and a given location in the FOA. The role of the selection network in the model is to gate activity in the contents network, enabling activity to be

passed on from some but not other units in the contents network. This 'selects' which mapping is implemented from the retina to the FOA. Through bottom-up activation in the Contents and Selection networks, a given stimulus will be mapped through to the FOA. However, in addition to this bottom-up form of operation, activation can also be conveyed in a top-down manner, from the Knowledge network down to the Selection net – essentially biasing activity in the Selection net to favour known over unknown stimuli. Finally, once a stimulus has been selected in the FOA and identified in the Knowledge network, SAIM utilizes a form of 'inhibition of return' (Posner & Cohen, 1984), in order to allow new (unselected) stimuli to be identified. In this process, the representations of an identified stimulus are inhibited, including the positions where the stimulus fell in a 'map of locations' (see Figure 6). This in turn allows previously un-selected items to then win the competition for selection, so that SAIM's 'attention' moves from one object to the next.

Figure 6 about here

The heart of SAIM's ability to select between multiple inputs, so that one object is identified at a time, is the Selection network. This network is shown in simplified form in Figure 7, where we depict a one-dimensional input from the retina being mapped, through the Selection network, to a one-dimensional representation of units in the FOA. Within the Selection network, each unit along a row corresponds to a different (but neighbouring) location in the visual field, and the units along one row all map into a single location in the FOA. Here units along a row will compete to control the mapping from the visual field into one position in the FOA. In contrast, each unit in a column of the Selection network corresponds to a different location in the FOA, and the units along one column all correspond to a single location in the visual field. Units along a column compete to control the mapping from one position in the visual field into the FOA. These competitive interactions were based on inhibitory connections between neighbouring units along each column and row. In addition, units in the Selection network are mutually excitatory if they support mappings from locally neighbouring units in the visual field into locally neighbouring units in the FOA. In Figure 6 this is illustrated by excitatory connections between units that lie along the diagonals of the Selection

network. These local connections, then, can be thought of as embodying various constraints about how spatial mapping should operate – for example, that one unit in the visual field should not map into more than one location in the FOA, and that neighbouring units in the field should map into neighbouring units in the FOA. This is similar to the way that certain computational constraints were built-into the connectionist model of stereopsis proposed in the classic model of Marr and Poggio (1976). When two objects are presented on SAIM's retina local units support one another, but more distant units are not mutually supportive and instead set-up competition to control the mapping from different parts of the field into the FOA. Objects may 'win' this competition either by having stronger bottom-up activation (e.g., if one object is larger than others, or if it has more elements packed around its centre of gravity (since such units form mutually-supportive alliances), or by receiving stronger top-down support from the Knowledge network (based either on an expectation formed before stimuli are presented or based on partial activation of the Knowledge network before selection has been completed).

Figure 7 about here

These points are illustrated in Figure 8, where we show the activity in SAIM when two objects are presented (in this case a + and a 2). Figure 8 presents activity at different time intervals in the FOA, as well as the activity that builds up (and is inhibited after recognition) in the Knowledge network. In this simulation the + is first selected in the FOA, and activation in its template increases to threshold level – the + is both attended and recognized. In this example, the + is attended first because it enjoys greater bottom-up support from the pixels surrounding its centre of gravity. Following this, representations for the + are inhibited. When processing continues, there is then a competitive advantage for the 2 over the +, so that the 2 then comes to be attended and identified. The stimuli are processed in parallel, but there is selection on one object at a time for the response.

Figure 8 about here

Within the framework presented by SAIM, unilateral neglect can come about by lesioning the Selection network. Heinke and Humphreys examined two different forms of lesioning. One form, which they termed a 'vertical lesion', affected units responding to input coming from one side of the visual field (e.g., all the units on the left side of the connection matrix presented in Figure 7). This meant that activity in these units suffered a competitive advantage relative to units in the Selection network responding to input present in the opposite (ipsilesional) side of the visual field. When two objects were presented, the ipsilesional object tended to be attended first, even if the bottom-up input based on the shape alone would favour the stimulus presented in the contralesional field. This is shown in Figure 9a, which illustrates performance of the model after a 'vertical lesion' affecting the left side of the Selection network. Although SAIM normally has a bottom-up preference for the + over the 2 (Figure 8), after lesioning, the model selects the 2 first if this item falls in the ipsilesional field and the + in the contralesional field. Furthermore, even after selecting the 2, the model has difficulty in selecting the + - partly because any spatial distortion in mapping the 2 into the FOA tends to make it difficult to inhibit the stimulus, and partly because the competitive advantage for the + is difficult to suppress. Hence only the 2 is attended in this example. This is also not simply a neglect of the contralesional part of the visual field, since SAIM does select and identify the contralesional + when it is the only object present (Figure 9b). There is a relative neglect of the more contralesional of two separate objects: there is neglect between separate objects.

Figure 9 about here

A different pattern of impairment was evident, however, if a 'horizontal' lesion was performed, for example affecting the units in the top-most rows on the Selection network. Units in each row of the Selection network map are projected to from across the visual field, but they map into one unit in the FOA. Thus a horizontal lesion affects the mapping into one part of the FOA and it does this in a translation-invariant manner. This is demonstrated in Figure 10. In Figure 10 SAIM is presented with two stimuli after a 'horizontal lesion' has been imposed in order to disrupt access into the left side of the FOA. In this case, the model is able to map both stimuli into the FOA, but in each case there is poor representation of the left most pixel. Here there is a form of sequential neglect, for the left parts of both stimuli, affecting even stimuli in the ipsilesional visual field (cf. Gainotti et al., 1986). That is, a horizontal lesion generates a form of 'within-object neglect'.

Figure 10 about here

These simulations illustrate that SAIM can produce the dissociating pattern of deficits found in neglect patients, and this pattern emerges as a natural consequence of the way that stimuli are mapped into a FOA to achieve translation-invariant object recognition. These simulations are suggestive that dissociations between within- and between-object neglect are not arbitrary, but rather they can be expected as a function of the particular brain lesion affecting a patient. Indeed, by combining the horizontal and vertical lesions in a single simulation, Heinke and Humphreys (2003) were able to demonstrate that there could be left neglect of within-object representations, along with right neglect of between-object representations, in a single simulation. This is the pattern observed in patient JR (Humphreys & Riddoch, 1994, 1995).

SAIM, then, shows the ability of explicit computational models to capture rich patterns of data, providing order out of the seemingly chaos of empirical results. As we have noted with the simulations of NAM, though, the ability of a model to accommodate a set of results does not prove that the model is correct. Nevertheless, it should be noted that the dissociations between within-object and between-object neglect are not at all easy to simulate in other models that do not employ SAIM's architecture. For instance, the MORSEL model proposed by Michael Mozer and colleagues (e.g., Mozer, 1991; Mozer, Halligan & Marshall, 1997) uses a two-dimensional (spatial) 'attentional module' that modulates activity at early stages of visual processing. Units in the attentional module have a one-to-one correspondence to units on the retina. Damage to one side of this attentional module produces difficulty in perceiving stimuli presented on the corresponding part of space, and this can be exacerbated when there are also objects present on the ipsilesional side, which generate lateral inhibition within the attentional module. As a consequence, visual neglect can be simulated. However, it is difficult for such a model to capture patterns of performance found in patients with within-object neglect. For instance, following right hemisphere damage such patients can show better identification of the right-most features of an object present in the left (contralesional) visual field, relative to the left-most features of an object present in the right (ipsilesional) field. There is then also poor identification of the left-most features in the left side object (see Humphreys & Heinke, 1998, for data). We suggest that this is because there is a problem mapping features into one side of a spatial FOA; these results emerge naturally in SAIM's framework (see Figure 10). In MORSEL, however, a form of complex lesioning would need to be imposed on the attentional module, so that local deficits affect

the left-most features of each object, but not right-side features falling in between. This is unlikely, and could not then explain why report of features at the apparently impaired locations would be good if the right sides of objects fell there.

Due to its use of top-down activity, SAIM is also able to simulate evidence showing that there is reduced neglect of known relative to unknown objects (Humphreys & Riddoch, 1994, 1995). In Figure 11a we provide an illustration of SAIM's performance when it is given separate templates for each of two letters (I and T) and it is lesioned to produce left-side neglect. When presented with two stimuli, there is neglect of the left-most object. In this case, the T was selected first, and then re-selected as this letter continued to win the competition for selection over the I. In Figure 11b we demonstrate performance in the model when a third template is added, corresponding to the word IT. Now the model is able to recovery both of the letters present, even though it has templates for each individual letter still, and so could be biased to identity the T and not the whole word. Top-down activation, from the word template, helps to bias attention to cover both of the letters present, so that both can be attended. In this case, though there is a deficit at a stage that produces input into the Knowledge network, interactivity between the Knowledge and Selection network affects activity at the earlier stage. The model has structural modularity, and so a lesion can be selectively imposed within the Selection network, but its operation is non-modular and interactive. We suggest that this property of neural network-like models is useful for capturing some of the complexity of neuropsychological data.

Figure 11 about here

Other approaches to model ling.

Although we have presented a case for using biologically-inspired models to simulate both normal and neuropsychological data, other forms of modeling can also play a useful part in helping our understanding of neurological impairments. One example is the formal mathematical model of attention provided by Claus Bundesen's 'Theory of Visual Attention' (TVA). TVA proposes that visual selection is based on parallel competition between stimuli to map onto templates that are then represented in visual short-term memory (VSTM), making the stimuli available for report. This competition is said to be influenced by a number of parameters including the strength of the sensory signal, the speed of processing, the attentional weight that may be applied to stimuli as a function of their relevance to the task, and the capacity of VSTM. Duncan et al. (1999, 2003) and Habekost and Bundesen (2003) have applied fits of the parameters in TVA to the performance of brain-lesioned patients when asked to report either all or a sub-section of letters present in multi-letter displays (on whole- and partial-report tasks; cf. Sperling, 1960, 1967). The parameters were specific to whether stimuli appeared in the ipsi- or contralesional fields. In the Duncan et al. study, several interesting results emerged. For example, they found that patients with unilateral parietal damage had altered parameters for the rate of stimulus encoding and for VSTM capacity for stimuli in both visual fields, not just for items appearing on the contralesional side. This provides some explanation for the reports of patients not only showing contralesional neglect but also neglect of ipsilesional items when their attention is drawn to the contralesional side (e.g., Robertson, 1989). Further insights using this approach have been gained into understanding the neuropsychological disorder of 'simultanagnosia', in which patients seem impaired at

responding when multiple stimuli are presented simultaneously to them (e.g., Kinsbourne & Warrington, 1962). The factors that generate this disorder remain poorly understood, with arguments ranging from impaired integration of information in VSTM (Coslett & Saffran, 1991) to impaired application of an attentional spotlight to space (Treisman, 1998). Analyzing whole and partial-report performance in terms of TVA, Duncan et al. (2003) showed that the major parameter change was in the 'rate of processing' parameter, and they suggested that drastic slowing of processing would lead to many of the phenomena associated with the disorder, in which only the most dominant visual elements are reported.

The approach taken in applying TVA to neuropsychological data has both similarities and differences in relation to the attempt to model performance using artificial neural networks. In both instances, investigators aim to take a well-specified model that provides a fit to normal data, and then to use parameter changes after a brain lesion to account for a disorder. With TVA, the investigator uses the data to derive the parameters (and the parameter change, when present). In neural network modelling, the experimenter typically changes the parameters (e.g., adding noise to an activation function or removing units) and then assesses the effects on output. The neural network approach offers the possibility of emergent behaviours that were not specified in the original parameter setting, which can be of interest in linking to human disorders, but there is the inherent problem of finding the appropriate parameter change. This last problem is by-passed in mathematical modelling.

Learning and adding extra biological constraints.

We argue that models such as NAM and SAIM aid our understanding of disorders that are difficult to account for in terms of standard 'box and arrow' models, where strict modularity is assumed. It is also the case that both models are examples of a class of neural network in which both the representations and the parameters within the model are set by the experimenter. One of the attractions of neural network modelling, however, is that networks can learn internal representations that are required to transform a given input into an output. Especially where learning takes place within a quasi-regular environment, where there is some systematicity in the relations between inputs and output, models in which learning is incorporated can generate interesting emergent properties, which can in turn provide insights into neuropsychological disorders. One example here would be Plaut et al.'s (1996) model of word naming. In one version of this model, the experimenters used dual inputs to an output system where units represented the phonological properties of words. One input was provided by units representing orthographic properties of stimuli. The second input corresponded to a semantic representation of a word. The contribution of the semantic input was fixed over time, but that of the orthographic input varied and was subject to learning, in which the connections between the orthographic and phonological units were altered using a form of back propagation (Hinton, 1989). Relative to when learning took place without the semantic input, Plaut et al. found that, in the 'dual route' version, the 'route' mapping orthography onto phonology became more specialised for regular spelling-sound correspondences, with the naming of irregular word more dependent on the semantic route. Lesioning the semantic route then gave rise to a strong pattern of surface dyslexia, in which both regular words and nonwords were produced correctly (where the nonwords

had highly regular spelling-sound correspondences) but irregular words were selectively impaired. This simulates patterns of 'pure' surface dyslexic reading (Shallice et al., 1983) that have been difficult to simulate in 'single route' models incorporating only a single pathway between orthography and phonology (e.g., Patterson, Seidenberg & McClelland, 1989). In Plaut et al. (1996), there was a 'division of labour' between the semantic and orthographic routes for reading as learning took place, so that a more 'regularised' representation was developed than when a single route was used. Clearly it is of interest to examine how learning may interact with structural constraints in an artificial neural network, to generate forms of representation that are 'tuned' in different ways.

A further example of using structural constraints to generate emergent properties in networks comes from studies where topographical biases have been incorporated into the learning process. Plaut (2002) used a distributed 'semantic' network with a two-dimensional topology to learn mappings between a variety of input stimuli and output tasks (vision and touch as input, naming and gesturing as output; see Figure 12). The model had a topological bias, in which short connections were favoured over long connections (the magnitudes of the changes made to connect weights during learning were greater for short connections than for long connections; see also Jacobs & Jordan, 1992). With this bias, units in the semantic network that were close to a particular input or output modality were strongly weighted to achieve a particular (modality-specific) mapping, whilst units that were more distant from the different input and output modalities played a more 'multi-modal' role, and were not differentially involved in particularly input-output mappings. It followed that lesions to the network generated different patterns of performance according to which units were affected. For example, a

pattern of optic aphasia arose (i.e., impaired naming but not gesturing to visually presented objects) from damage to connections to semantic units specialised for mapping visual input onto names (e.g., units in the lower left region of the two-dimensional semantic space in Figure 12). Problems in gesturing to visual input (i.e., the pattern of visual apraxia) were apparent after damage to connections to units specialised for mapping between visual input and gestural output (i.e., units in the upper left region of the semantic space shown in Figure 12). This, then, produces a pattern of performance that is not dissimilar from the 'dual route' NAM, with functional specialisation in the distributed network developing because of the topological constraints (indeed, the same patterns of performance did not arise when the constraints were eliminated, in control simulations). One difference between NAM and the structured, distributed semantic model of Plaut (2002) is that, in NAM, visually presented objects must be named through the semantic system – there are no direct connections from visual structural representations of objects to names. In contrast, in Plaut's model, some units can specialise for mapping from vision to names, and indeed lesioning of these units tended to produce optic aphasia. However, any evidence for 'direct visual naming of objects' in the neuropsychological literature is weak (e.g., instances where patients name an object but cannot retrieve other semantic knowledge; see Hodges & Greene, 1998; Humphreys & Forde, in press). Thus there is not independent support for the argument that there can be neural specialisation for 'direct' mapping from vision to names.

Conclusions

Neuropsychological research can provide striking insights into the nature of mental processes, dissecting apart processes that are functionally independent of one another. There are instances, however, where the assumptions that have historically been used to guide theorising have difficulty in accounting for some neuropsychological disorders (e.g., visual apraxia). We have argued that the development and use of explicit computational (and mathematical) models can be beneficial here, since such models can capture forms of interactivity between processing modules, that turn out to be important for understanding human performance. Such models can also provide a framework for understanding the relations between different neuropsychological disorders, along with generating formal accounts of how a brain lesion can selectively affect different processing parameters. There are examples where such formal accounts suggest new insights into the nature of the disorder. In addition to this, models that employ learning, and that incorporate additional biological constraints into their learning functions, can generate emergent properties that can aid our understanding of the relations between brain structure and cognitive function. We conclude that computational studies are an important 'tool' for cognitive neuroscientific analysis of brain and mind.

Acknowledgements.

This work was supported by grants from the BBSRC, EPSRC and MRC to Glyn Humphreys and Dietmar Heinke, and from the British Academy to Eun Young Yoon.

References

Beauvois, M-F. & Derouesné, J. (1979) Phonological alexia: Three dissociations. Journal of Neurology, Neurosurgery and Psychiatry, 42, 1115-1124.

Bub, D., Cancelliere, A. & Kertesz, A. (1985) Whole-word and analytic translation of spelling-to-sound in a non-semantic reader. In K.E. Patterson, M. Coltheart & J.C. Marshall (Eds.), Surface dyslexia. London: Erlbaum.

Caramazza, A. (1986) On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. Brain and Cognition, 5, 41-66.

Coltheart, M. (1984) Editorial. Cognitive Neuropsychology, 1, 1-8.

Coltheart, M. (1985) Cognitive neuropsychology and the study of reading. In M.I. Posner & O.S.M. Marin (Eds.), Attention and performance XI. Hillsdale, N.J.: Erlbaum.

Coltheart, M., Curtis, B., Atkins, P. & Haller, M. (1993) Models of reading aloud: Dual-route and parallel distributed processing approaches. Psychological Review, 100, 589-608.

Coslett, H.B. & Saffran, E. (1991) Simultanagnosia: To see but not two see. Brain, 114, 1523-1545.

DeRenzi, E., Faglioni, P. & Sorgato, P. (1982) Modality specific and supramodal mechanisms of apraxia. Brain, 105, 301-312.

Duncan, J., Bundesen, C., Olson, A., Humphreys, G.W., Chavda, S. & Shibuya, H. (1999). Systematic analysis of deficits in visual attention. Journal of Experimental Psychology: General, 128, 1-29.

Duncan, J., Bundesen, C., Olson, A., Humphreys, G.W., Ward, R., van Raamsdonk, M., Rorden, C. & Chavda, S. (2003) Attentional functions in dorsal and ventral simultanagnosia. Cognitive Neuropsychology, 20, 675-702.

Freund, D.C. (1889) Uber optische Aphasie und Seelenblindheit. Archiv für Psychiatrie und Nervenkrankheiten, 20, 276-297.

Funnell, E. (1983) Phonological processing in reading: New evidence from acquired dyslexia. British Journal of Psychology, 74, 159-180.

Gainotti, G., D'Erme, P., Monteleone, D. & Silveri, M.C. (1986) Mechanisms of unilateral spatial neglect in relation to laterality of cerebral lesion. Brain, 109, 599-612.

Habekost, T. & Bundesen, C. (2003) Patient assessment based on a theory of visual attention (TVA): Subtle deficits after a right frontal-subcortical lesion. Neuropsychologia, 41, 1171-1188.

Halligan, P. W., Fink, G. R., Marshall, J. C., & Vallar, G. (2003). Spatial cognition: evidence from visual neglect. Trends in Cognitive Sciences, 7, 125-133.

Heilman, K.M., Watson, R.T. & Valenstein, E. (1985) Neglect and related disorders. In K.M. Heilman & E. Valenstein (Eds.), Clinical neuropsychology. Oxford: Oxford University Press.

Heinke, D. & Humphreys, G.W. (2003) Attention, spatial representation and visual neglect: Simulating emergent attentional processes in the Selective Attention for Identification Model (SAIM). Psychological Review, 110, 29-87.

Hillis, A.E. & Caramazza, A. (1995) Cognitive and neural mechanisms underlying visual and semantic processing: Implications from 'optic aphasia'. Journal of Cognitive Neuroscience, 7, 457-478.

Hinton, G.E. (1989) Connectionist learning procedures. Artificial Intelligence, 40, 185-234.

Hodges, J.J., Bozeat, S., Lambon Ralph, M.A., Patterson, K. & Spatt, J. (2000) The role of conceptual knowledge in object use: Evidence from semantic dementia. Brain, 123, 1913-1925.

Hodges, J.R. & Greene, J.D.W. (1998) Knowing about people and nursing them: Can Alzheimer's Disease patients do one without the other? Quarterly Journal of Experimental Psychology, 51A, 121-134.

Humphreys, G.W. & Heinke, D. (1998). Spatial representation and selection in the brain: Neuropsychological and computational constraints. Visual Cognition, 5, 9-47.

Humphreys, G.W. & Riddoch, M.J. (1994) Attention to within-object and between-object spatial representations: Multiple sites for visual selection. Cognitive Neuropsychology, 11, 207-242.

Humphreys, G.W. & Riddoch, M.J. (1995). Separate coding of space within and between perceptual objects: Evidence from unilateral visual neglect. Cognitive Neuropsychology, 12, 283-312.

Humphreys, G.W. & Riddoch, M.J. (2003) From vision to action, and action to vision: A convergent route approach to vision, action and attention. In D. Irwin & B. Ross (Eds.), The psychology of learning and motivation: Visual cognition. V. 42. New York: Academic Press.

Jacobs, R.A. & Jordan, M.I. (1992) Computational consequences of a bias toward short connections. Journal of Cognitive Neuroscience, 4, 323-336.

Kinsbourne, M. & Warrington, E.K. (1962) A disorder of simultaneous form perception. Brain, 85, 461-486.

Lhermitte, F. & Beauvois, M-F. (1973) A visual-speech disconnexion syndrome : Report of a case with optic-aphasia, agnosic alexia and colour agnosia. Brain, 96, 695-714.

Lichtheim, L. (1885) On aphasia. Brain, 7, 433-484.

Marr, D. & Poggio, T. (1976) Cooperative computation of stereo disparity. Science, 194, 283-287.

McClelland, J.L. & Rumelhart, D.E. (Eds.) (1986) Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2. Cambridge, Mass.: MIT Press.

Mozer, M.C. (1991) The perception of multiple objects: A connectionist approach. Cambridge, Mass.: MIT Press.

Mozer, M.C., Halligan, P.W. & Marshall, J.C. (1997) The end of the line for a brain-damaged model of unilateral neglect. Journal of Cognitive Neuroscience, 9, 171-190.

Patterson, K.E., Seidenberg, M.S. & McClelland, J.L. (1989) Connections and disconnections: Acquired dyslexia in a computational model of reading processes. In R.G.M. Morris (Ed.), Parallel distributed processing: Implications for psychology and neuroscience. London: Oxford University Press.

Pilgrim, E. & Humphreys, G.W. (1991) Impairment of action to visual objects in a case of ideomotor apraxia. Cognitive Neuropsychology, 8, 459-473.

Pizzamiglio, L. et al. (1989) Visual neglect for far and near extra-personal space in humans. Cortex, 25, 471-477.

Pitzalis, S. et al. (2001) Influence of the radial and vertical dimensions on lateral neglect. Experimental Brain Research, 136, 281-294.

Plaut, D.C. (2002) Graded modality-specific specialisation in semantics: A computational account of optic aphasia. Cognitive Neuropsychology, 19, 603-639.

Plaut, D.C., McClelland, J.L., Seidenberg, M.S. & Patterson, K.E. (1996) Understanding normal and impaired word reading: Computational principles in quasi-regular domains. Psychological Review, 103, 56-115.

Posner, M.I. & Cohen, Y. (1984) Components of visual orienting. In H. Bouma & D.G. Bouwhuis (Eds.), Attention and performance X. pp. 531-556. Hillsdale, N.J.; Lawrence Erlbaum Associates.

Riddoch, M.J. & Humphreys, G.W. (1987) Visual object processing in optic aphasia: A case of semantic access agnosia. Cognitive Neuropsychology, 4, 131-185.

Riddoch, M.J., Humphreys, G.W. & Price, C.J. (1989) Routes to action: Evidence from apraxia. Cognitive Neuropsychology, 6, 437-454.

Robertson, I. (1989) Anomalies in the lateralisation omissions in unilateral left neglect: Implications for an attetional theory of neglect. Neuropsychologia, 27, 157-165.

Rumelhart, D.E. & McClelland, J.L. (Eds.) (1986) (1986) Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1. Cambridge, Mass.: MIT Press.

Rumiati, R.I. & Humphreys, G.W. (1998). Recognition by action: Dissociating visual and semantic routes to action in normal observers. Journal of Experimental Psychology: Human Perception and Performance, 24, 631-647.

Shallice, T., Warrington, E.K. & McCarthy, R.A. (1983) Reading without semantics. Quarterly Journal of Experimental Psychology, 35A, 111-138.

Sperling, G. (1960) The information available in brief visual presentations. Psychological Monographs, 74, (11, Whole No. 498)

Sperling, G. (1967) Successive approximations to a model for short-term memory. In A.F. Sanders (Ed.), Attention & performance I. Amsterdam: North Holland.

Treisman, A. (1998) Feature binding, attention and object perception. Philosophical Transactions of the Royal Society, 353, 1295-1306.

Yoon, E.Y., Heinke, D. & Humphreys, G.W. (2002) Modelling direct perceptual constraints on action selection: The Naming and Action Model (NAM). Visual Cognition, 9, 615-661.

Yoon, E.Y., Humphreys, G.W. & Riddoch, M.J. (in press) Action naming with impaired semantics: Neuropsychological evidence contrasting naming and reading for objects and verbs. Cognitive Neuropsychology.

Figure legends.

Figure 1. A 'dual route' model of naming English words. In this model there are distinct 'lexical' and 'non-lexical' routes for translating from spelling to sounds, each composed of a variety of processes (e.g., an orthographic lexicon, a semantic system, a phonological lexicon, within the lexical route). The assumption of modularity holds that one representation can be damaged without affecting the performance of indepednent representations. Thus damage to a non-lexical route should not affect the operation of the lexical route.

Figure 2. A 'dual route' model of action retrieval from visually presented objects, separating a direct visual and an indirect (semantically) mediated route to action. The model also proposes that access to object names from vision is semantically mediated (following Riddoch et al., 1989).

Figure 3. Architecture of the Naming and Action Model (NAM) (after Yoon et al., 2002), with inputs provided either by structural descriptions of objects or words. The mapping from each form of structural description into the Semantic system was based on a radial basis function network.

Figure 4. Data on action classification to objects and words, along with name retrieval to objects following simulated damage to the direct visual (V→ A) and indirect semantic routes to action (V→ S), in NAM (after Yoon et al., 2002). There is better action retrieval to objects after the V→S lesion than the V→A lesion, but better visual access to object names after the V→A than the V→S lesion. In both cases, action retrieval to words is spared.

Figure 5. Data on action classification and naming to a deadline in NAM (after Yoon et al., 2002). Here we present the proportion of error responses when retrieving actions to visually presented objects or to words. Relative to when actions were retrieved to words, there were more visual errors when retrieving actions to objects, and fewer semantic + semantic/visual errors (semantic and semantic+visual errors are classed together here because the two error types are difficult to distinguish).

Figure 6. The architecture of the Selective Attention for Identification Model (SAIM) (after Heinke & Humphreys, 2003).

Figure 7. Illustration of the Selection network in SAIM. Input from the visual field is depicted in terms of the bottom row of units, whilst units in the FOA are illustrated by the vertical row on the left. → indicates an excitatory connection; --o indicates an inhibitory connection.

Figure 8. Activity in SAIM when two stimuli (a + and a 2) are presented. There is sequential identification of the items, with a bottom-up advantage for the + (which is selected first, followed by the 2, after there is inhibition of return for the +).

Figure 9. (a) Activity in SAIM after a 'vertical' lesion has affected the left side of the Selection network, when two stimuli are presented (b) Activity in SAIM following the same lesion as in (a) when a single stimulus is presented in different positions of the field. In this last case, the + is always attended, although the time to be attended (time FOA) amd to be identified (time template) varies across the field (slower on the left).

Figure 10. An example of ‘within object’ neglect followiung a ‘horizontal’ lesion of the Selection network. Both stimuli are attended, but the left-most pixels are excluded from the FOA.

Figure 11. (a) Activity in SAIM after a ‘vertical’ lesion when the model has separate templates for the two letters presented. (b) Activity in SAIM when subject to the same lesion as in (a), but there is now an added template that incorporates both letters.

Figure 12. Illustration of the structured distributed semantic memory system proposed by Plaut (2002).

Figure 1

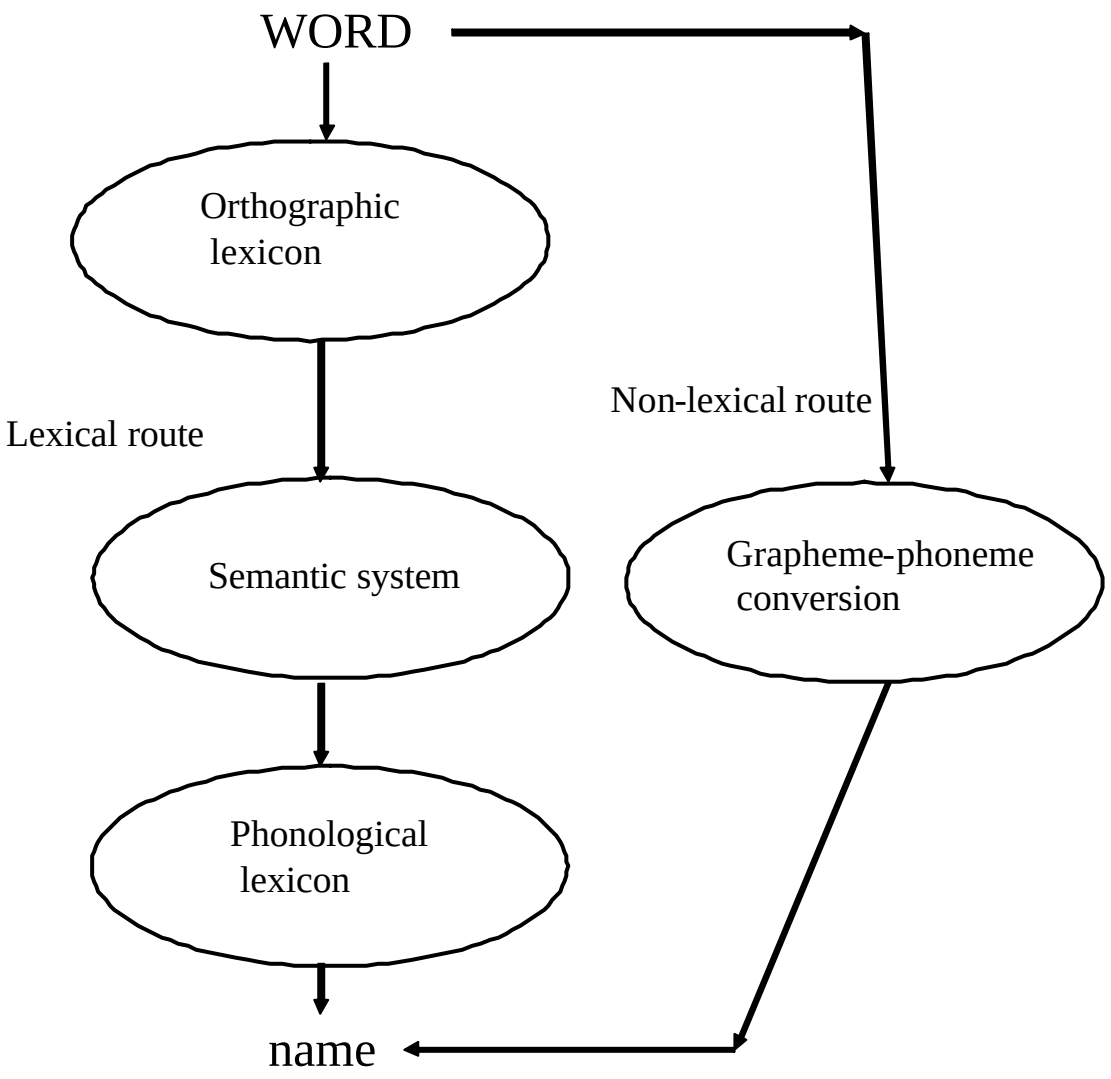
WORD
Orthographic lexicon
Lexical route
Non-lexical route
Semantic system
Grapheme-phoneme conversion
Phonological lexicon
name

Figure 2.

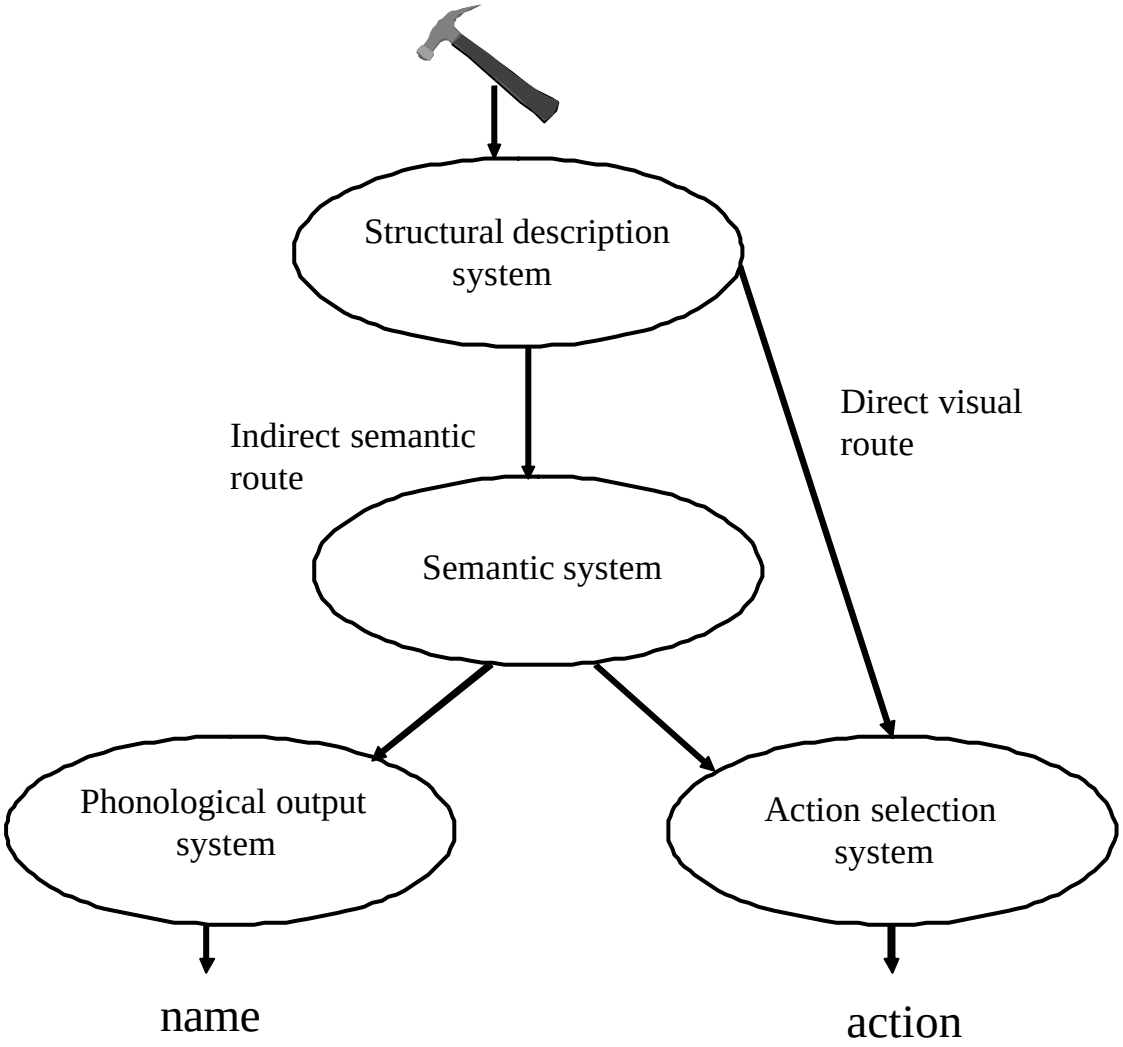
Structural description system
Indirect semantic route
Direct visual route
Semantic system
Phonological output system
Action selection system
name
action

Figure 3.

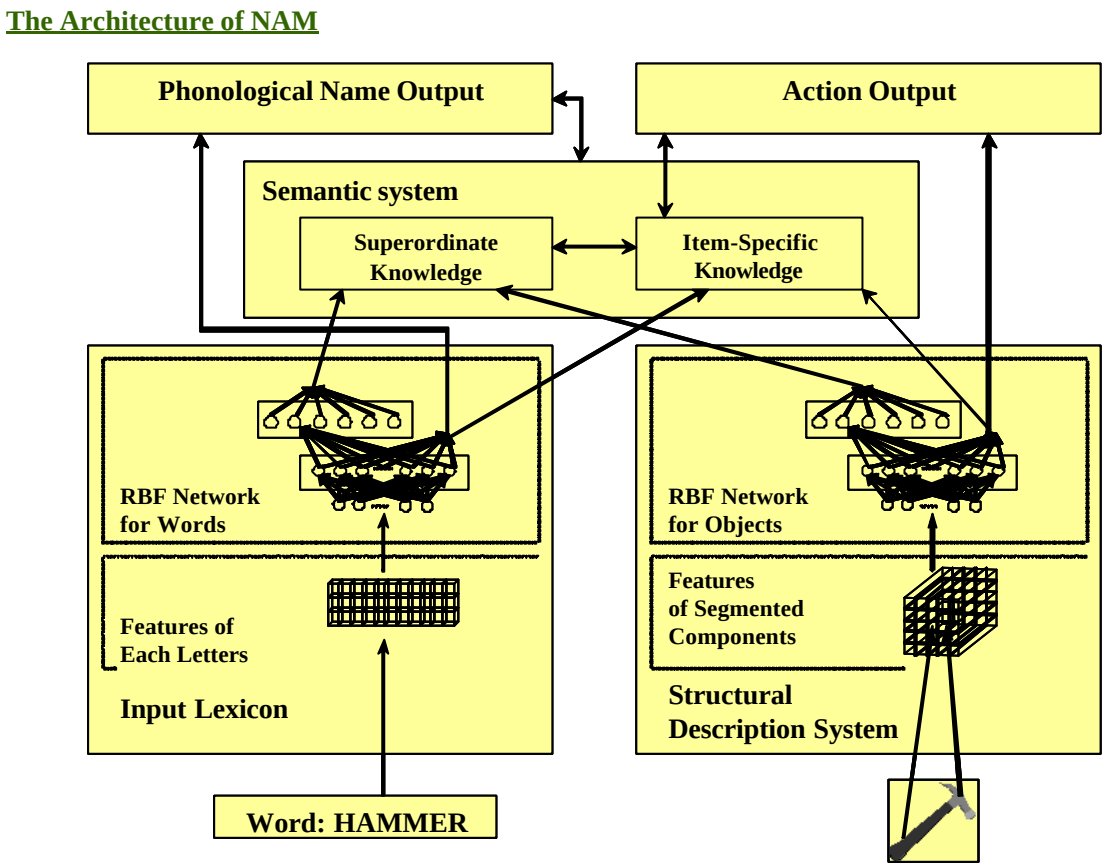
The Architecture of NAM
Phonological Name Output
Action Output
Semantic system
Superordinate Knowledge
Item-Specific Knowledge
RBF Network for Words
Features of Each Letters
Input Lexicon
RBF Network for Objects
Features of Segmented Components
Structural Description System
Word: HAMMER

Figure 4.

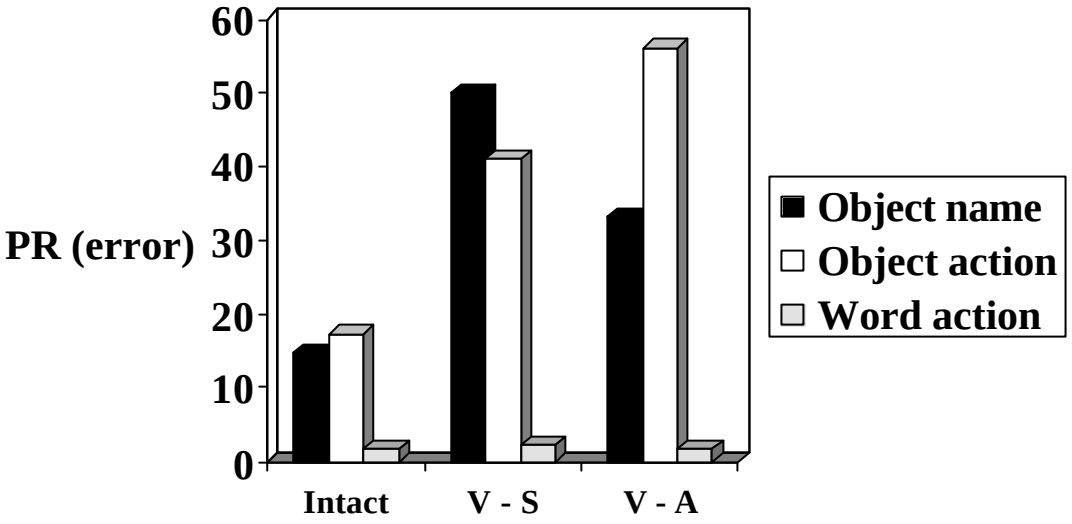
PR (error)
60
50
40
30
20
10
0
Intact
V - S
V - A
Object name
Object action
Word action

Figure 5.

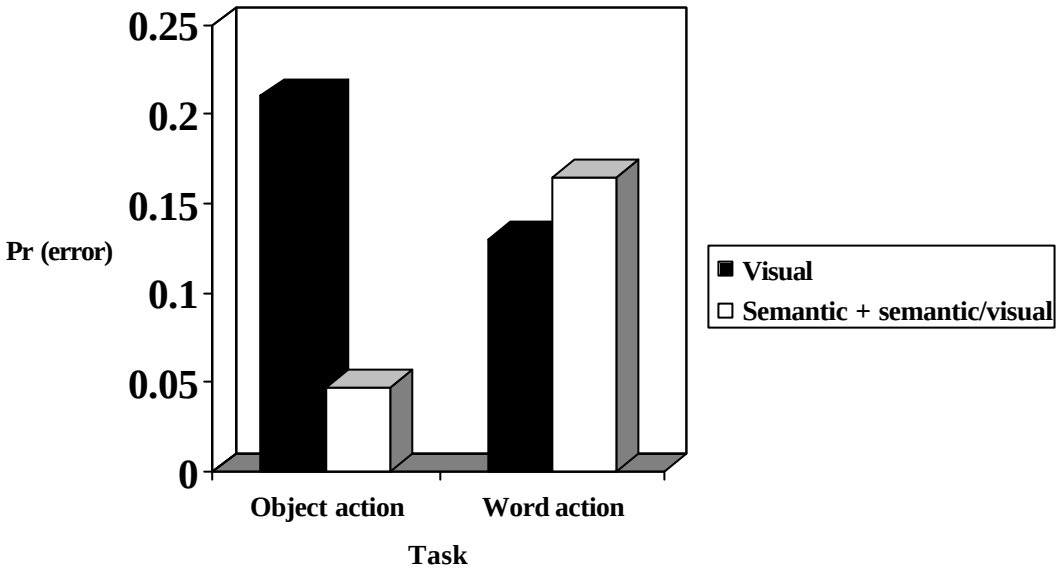
0.25
0.2
0.15
0.1
0.05
0
Pr (error)
Object action
Word action
Task
Visual
Semantic + semantic/visual

Figure 6

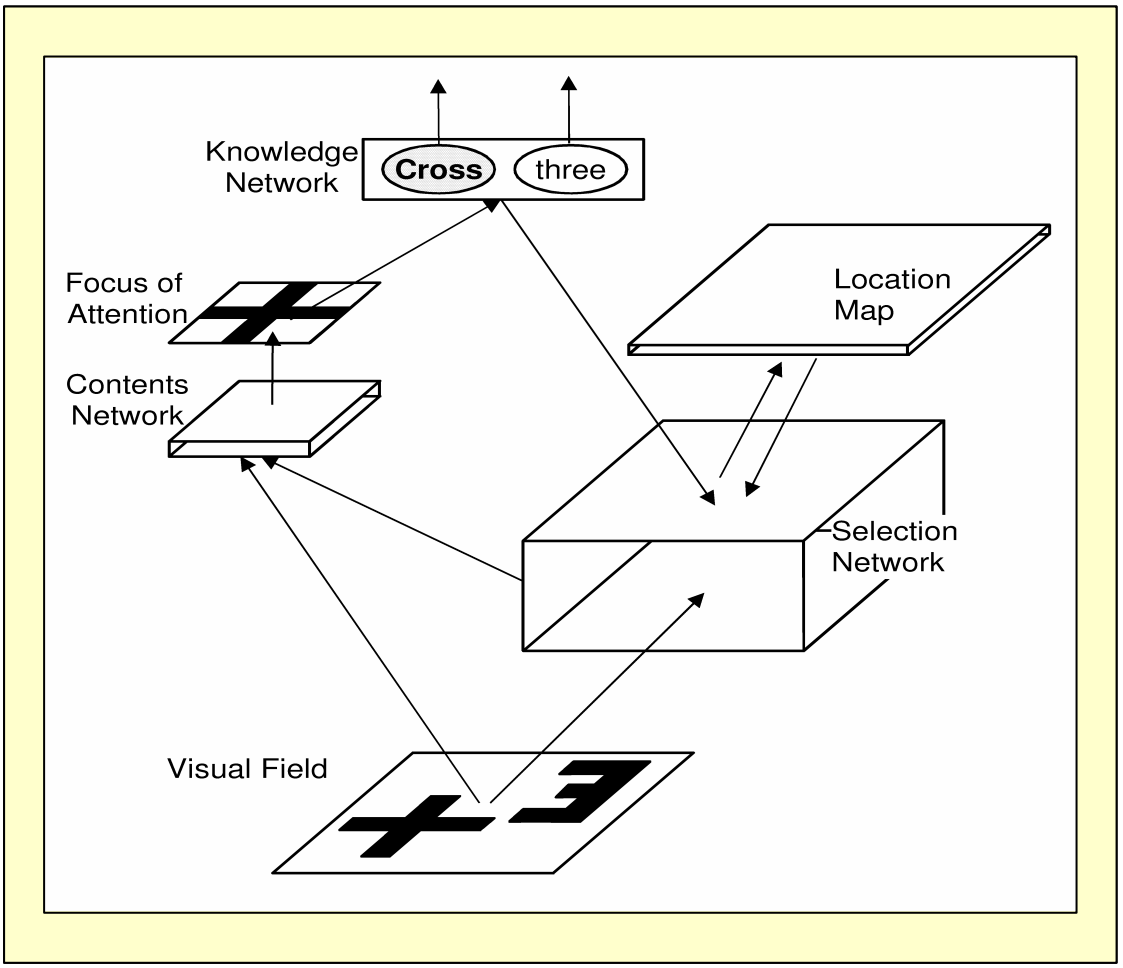
Knowledge Network
Cross
three
Focus of Attention
Location Map
Contents Network
Selection Network
Visual Field

Figure 7

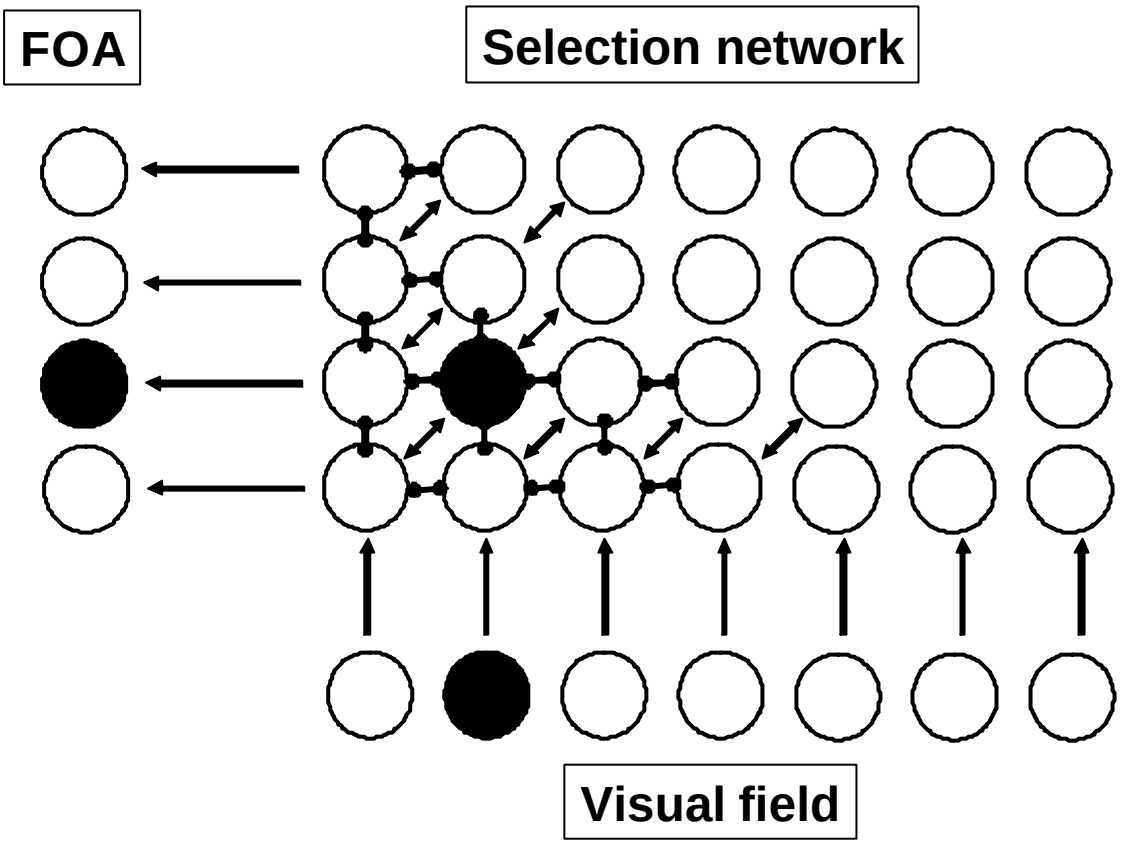

Figure 8

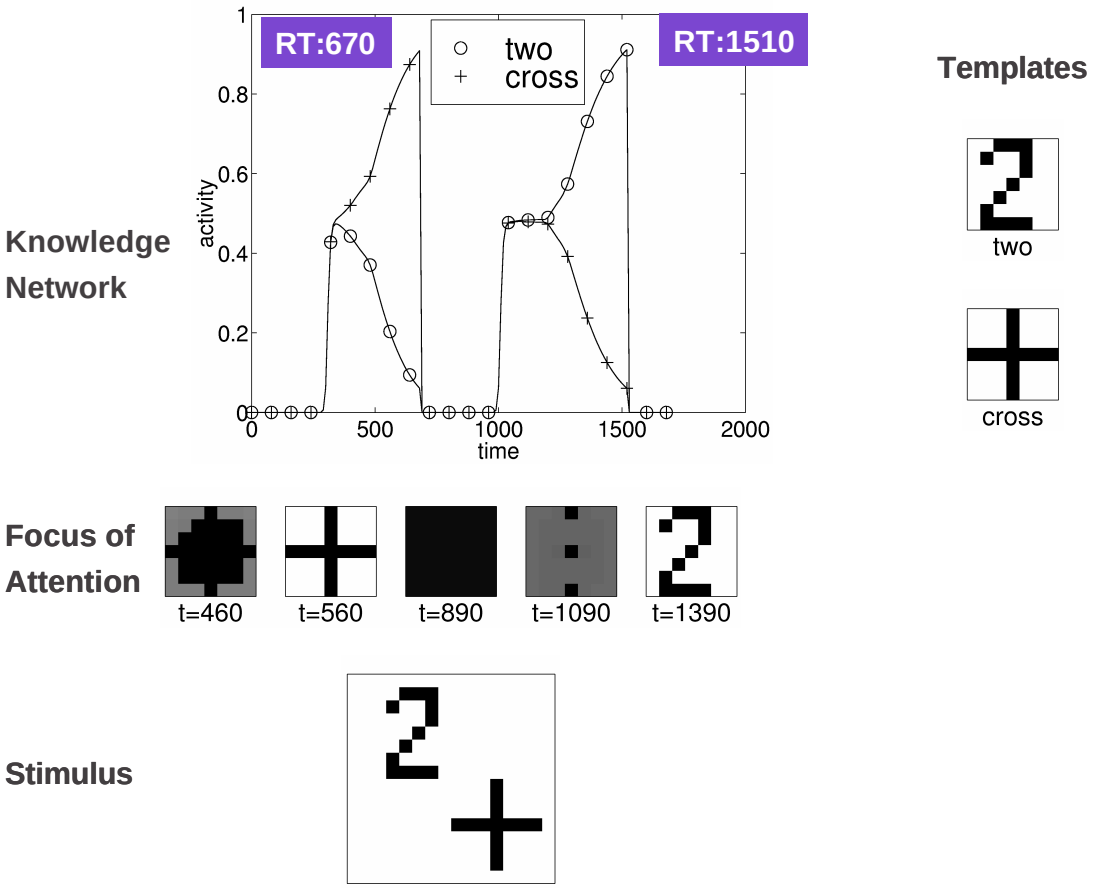
Knowledge Network
RT:670
RT:1510
two
cross
activity
time
0
500
1000
1500
2000
0
0.2
0.4
0.6
0.8
1
Templates
two
cross
Focus of Attention
t=460
t=560
t=890
t=1090
t=1390
Stimulus

Figure 9 (a)

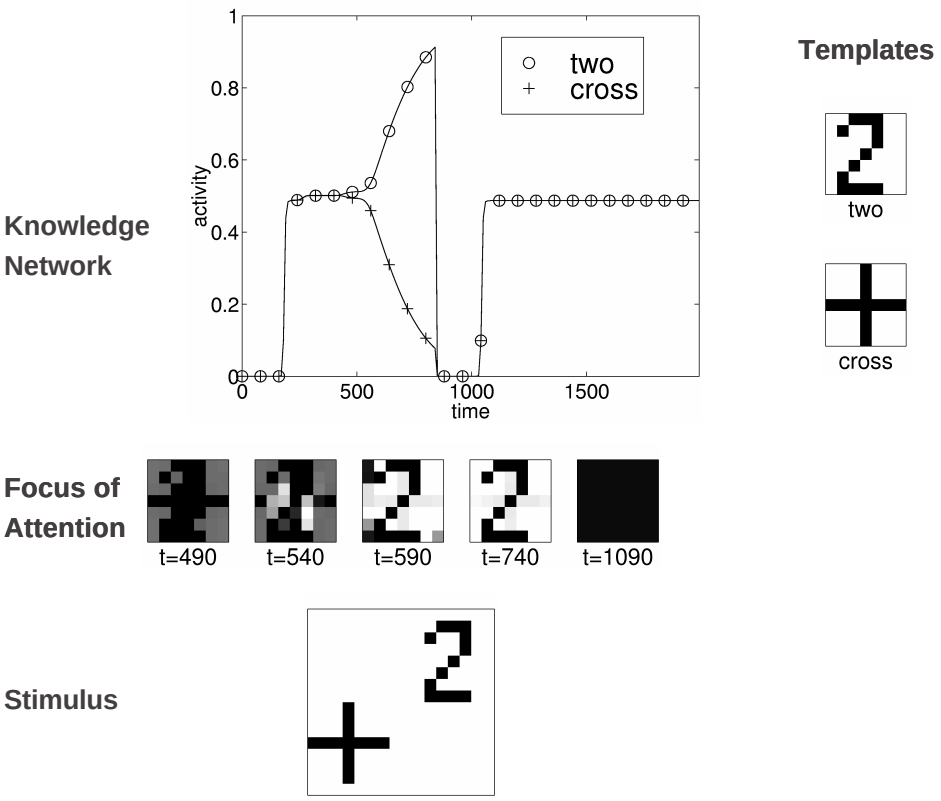
Knowledge Network
activity
time
two
cross
Templates
two
cross
Focus of Attention
t=490
t=540
t=590
t=740
t=1090
Stimulus

Figure 9 (b)

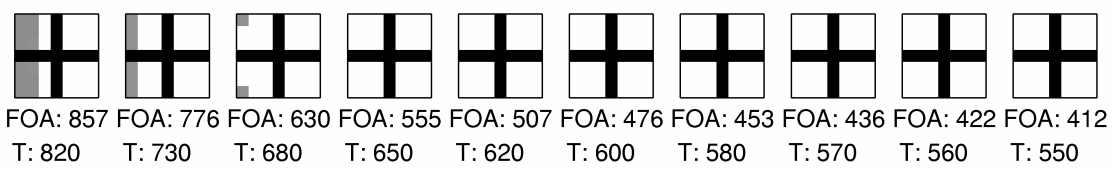

Focus of attention for different locations

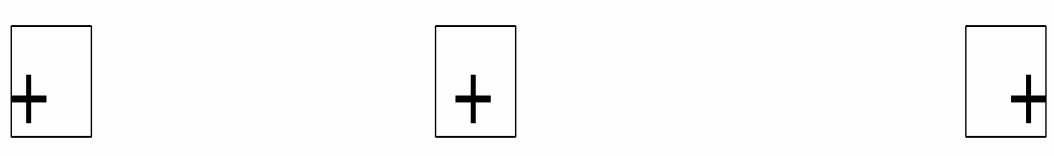

Stimuli: Cross at different locations

Figure 10

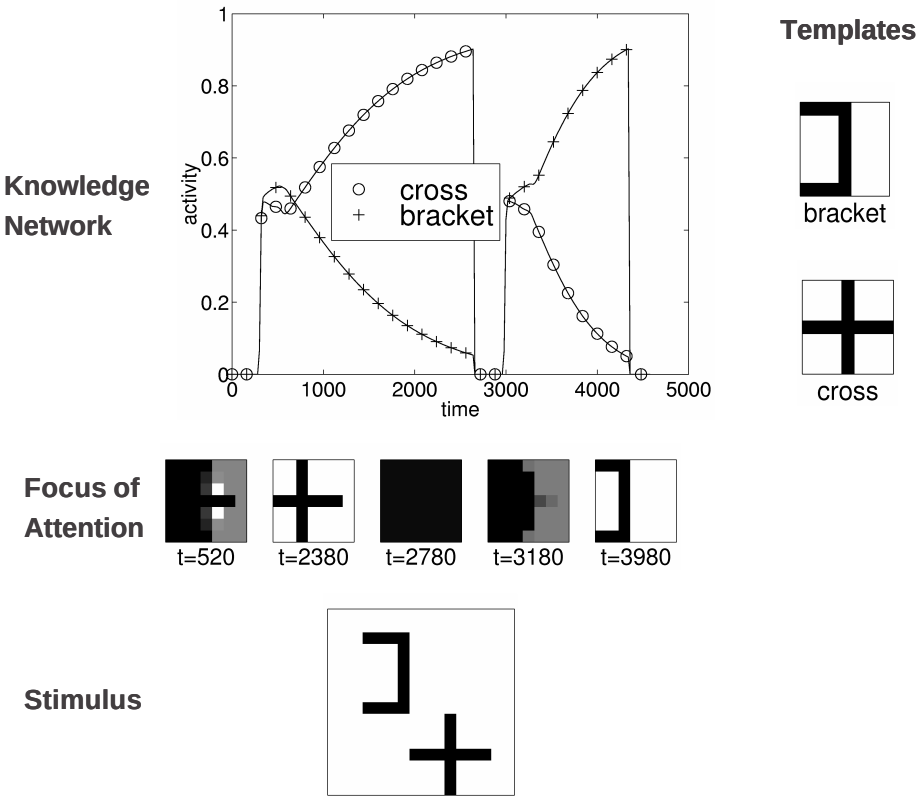
Templates
Knowledge Network
activity
time
cross
bracket
bracket
cross
Focus of Attention
t=520
t=2380
t=2780
t=3180
t=3980
Stimulus

Figure 11 (a)

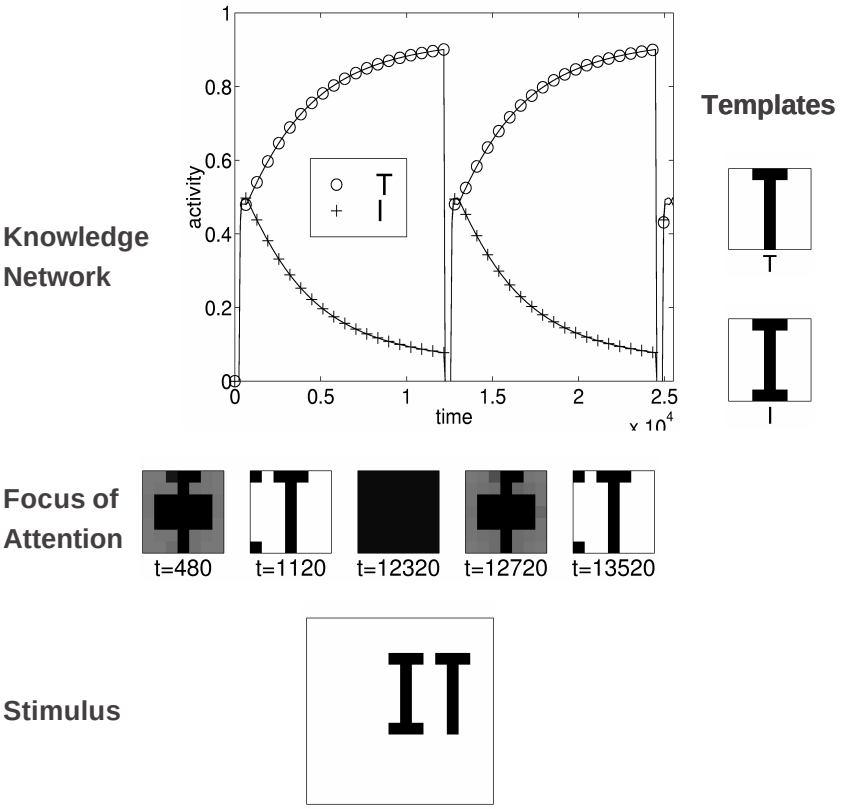
Knowledge Network
activity
time
x 10^4
T
I
Templates
T
I
Focus of Attention
t=480
t=1120
t=12320
t=12720
t=13520
Stimulus

Figure 11 (b)

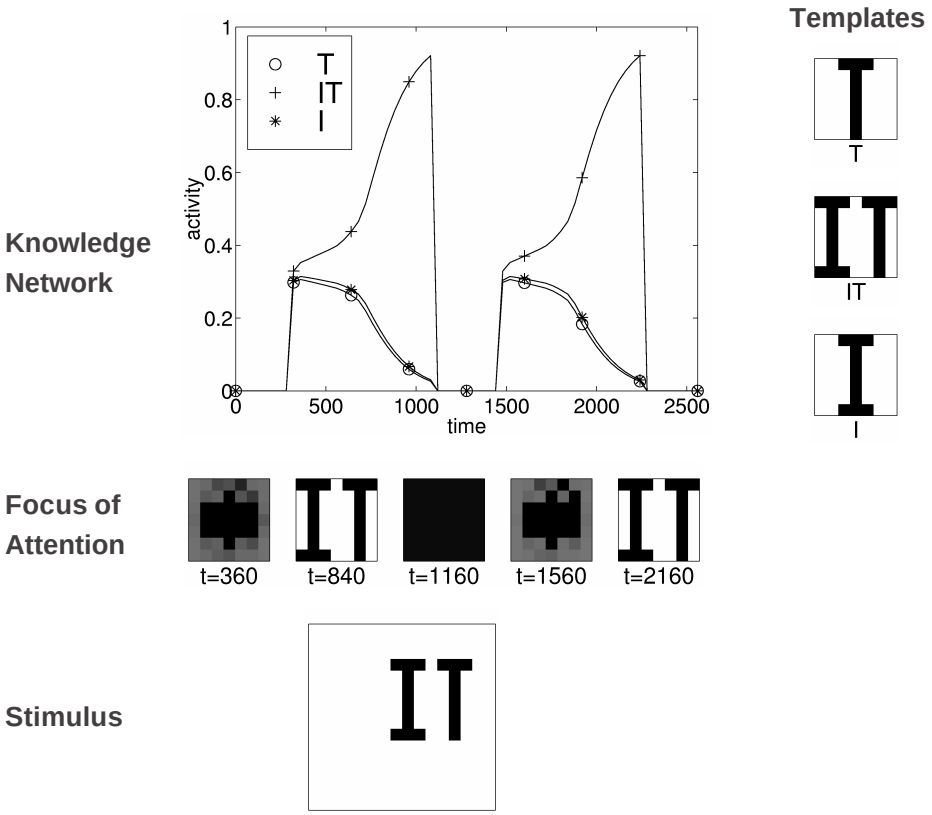
Templates
Knowledge Network
activity
time
T
IT
I
Focus of Attention
t=360
t=840
t=1160
t=1560
t=2160
Stimulus

Figure 12.